# TRPL+K: THICK-RESTART PRECONDITIONED LANCZOS+K METHOD FOR LARGE SYMMETRIC EIGENVALUE PROBLEMS* 

LINGFEI $\mathrm{WU}^{\dagger}$, FEI $\mathrm{XUE}^{\ddagger}$, AND ANDREAS STATHOPOULOS ${ }^{\dagger}$


#### Abstract

The Lanczos method is one of the standard approaches for computing a few eigenpairs of a large, sparse, symmetric matrix. It is typically used with restarting to avoid unbounded growth of memory and computational requirements. Thick-restart Lanczos is a popular restarted variant because of its simplicity and numerical robustness. However, convergence can be slow for highly clustered eigenvalues so more effective restarting techniques and the use of preconditioning is needed. In this paper, we present a thick-restart preconditioned Lanczos method, TRPL +K , that combines the power of locally optimal restarting $(+\mathrm{K})$ and preconditioning techniques with the efficiency of the thick-restart Lanczos method. TRPL+K employs an inner-outer scheme where the inner loop applies Lanczos on a preconditioned operator while the outer loop augments the resulting Lanczos subspace with certain vectors from the previous restart cycle to obtain eigenvector approximations with which it thick restarts the outer subspace. We first identify the differences from various relevant methods in the literature. Then, based on an optimization perspective, we show an asymptotic global quasi optimality of a simplified TRPL +K method compared to an unrestarted global optimal method. Finally, we present extensive experiments showing that TRPL+K either outperforms or matches other state-of-the-art eigenmethods in both matrix-vector multiplications and computational time.


Key words. symmetric eigenvalue problems, thick-restart, preconditioned Lanczos, global quasi optimality

AMS subject classifications. 65F15, 15A18
DOI. 10.1137/17M1157568

1. Introduction. The numerical solution of large sparse symmetric eigenvalue problems is one of the most computationally intensive tasks in applications ranging from structural engineering, quantum chromodynamics, material science, dynamical systems, machine learning, and data mining to numerical linear algebra $[10,31,39,4,45,48,47,46,15,13,44,49]$. Depending on the application, one may be interested in computing a few of the smallest or largest eigenpairs, or some eigenpairs in the interior of the spectrum. The challenge is that the size of the eigenproblems in these applications is routinely $O\left(10^{7}-10^{9}\right)$ [26]. Due to memory and computational constraints, iterative methods that rely on sparse matrix-vector products are the only practical solutions.

We are interested in finding the smallest eigenvalues and associated eigenvectors of the pencil $(A, B)$ when $A, B \in \mathbb{R}^{n \times n}$ are large, sparse symmetric matrices,

$$
\begin{equation*}
A v_{i}=\lambda_{i} B v_{i}, \quad i=1, \ldots, p, \quad p \ll n \tag{1.1}
\end{equation*}
$$

where $B$ is positive definite, $\left[v_{1}, \ldots, v_{p}\right] \in \mathbb{R}^{n \times p}$ is a $B$-orthonormal set of eigenvectors, and $\lambda_{1}, \ldots, \lambda_{p}$ are the corresponding eigenvalues of $(A, B)$. For simplicity, we first

[^0]discuss our method in the context of the standard eigenvalue problem (where $B$ is the identity matrix) but we extend it to the generalized problem later.

Variants of Krylov subspace methods have long been used to address large-scale eigenvalue problems $[19,27,28,29,16]$. The unrestarted versions of Lanczos and Arnoldi are considered optimal methods because they obtain their solutions over the entire set of matrix polynomials with degree up to the number of iterations. Yet, for difficult problems they require too many iterations to converge, resulting in impractical memory and computational demands. Researchers have sought to alleviate these problems with restarting and preconditioning.

Restarting interrupts the iteration, computes the current approximations, and uses them to start a new iteration. The pioneering implicitly restarted Arnoldi and Lanczos (IRL) methods perform this in a way so that the restarted vectors continue to form a Krylov subspace [33, 2]. The thick-restart Lanczos (TRLan) method [22, 43] is equivalent to IRL but simpler to use and numerically robust. However, when the desired eigenvalues are poorly separated from the rest of the spectrum, restarting causes further deterioration of convergence that thick restarting cannot fully recover. Locally optimal restarting is a technique that can result in near-optimal convergence when combined with thick restarting or block methods $[17,36,34]$. The technique is also known as +K restarting, which is more descriptive of the number of locally optimal restarting directions we keep. The resulting methods, however, do not form Krylov spaces and cannot use efficient constructing strategies such as the Lanczos three-term recurrence.

Near-optimal restarting techniques alone cannot address the slow convergence caused by poorly separated eigenvalues, a fact that often appears in real applications [50, 42]. Shift-invert transformations require matrix inversions and are typically expensive, and we do not consider them in this paper. Instead we focus on methods that use inexact inverses, or preconditioning, to accelerate convergence. The Davidson method and its extension, the generalized Davidson (GD) method, are prototypical preconditioned methods $[6,23,5,35]$. At every step they apply the preconditioner to the current eigenvalue residual to extend the search space in a similar fashion to Arnoldi but producing a non-Krylov space. The Jacobi-Davidson method (JD) [32] is a special case of GD where the preconditioner is performed with an appropriately projected preconditioned linear system solver. GD can also use thick and locally optimal restarting, a method called GD+K [36]. Note that this type of eigenvalue preconditioning can be considered as a step of an optimization method [9]. Such a view is followed by the locally optimal block preconditioned conjugate gradient method (LOBPCG) [17], which forgoes the subspace acceleration of GD+K for a block threeterm recurrence. The search spaces of the above methods are not Krylov spaces, which results in two disadvantages: expensive iteration costs (Rayleigh-Ritz (RR) projection at each inner step) and selective convergence to a particular eigenpair. The preconditioned Lanczos (PL) [24] and the inverse free preconditioned Krylov subspace (EigIFP) methods [11] build a Krylov space of the preconditioned matrix and thus avoid the expensive iteration costs. Although the preconditioned matrix has different eigenvectors than $(A, B)$, the methods invoke the RR projection of the original matrices onto a preconditioned search subspace and can converge to one eigenpair at a time.

In this paper, we propose a thick-restart preconditioned Lanczos+K method (TRPL +K ) to address the aforementioned problems. TRPL +K includes all three major building blocks: thick restarting, locally optimal restarting, and preconditioning. Unlike GD+K, however, it employs a Krylov inner iteration based on TRLan
to build the search space, thus avoiding the expensive RR procedure at every step and requiring about half the memory. It also differs from JD, since it stores the entire Krylov space. Alternatively, TRPL+K can be considered as an extension to EigIFP (or PL) with thick and locally optimal restarting, thus offering significant speedups. We provide a convergence analysis of a simplified version of TRPL +K from the perspective of optimization, showing an asymptotic global quasi optimality of the method compared to an ideal unrestarted global optimal method. This complements some limited earlier theoretical results on the convergence of the +K technique in LOBPCG and GD+K [17, 36]. Extensive experiments demonstrate that TRPL +K either outperforms or matches other state-of-the-art eigenmethods in both matrix-vector multiplication counts and computational time.

In section 2 we develop a background framework on which to compare existing and proposed methods. In section 3, we present our method for the standard eigenproblem. Section 4 considers a simplified version of the algorithm for the generalized eigenproblem and develops a convergence analysis. Section 5 compares the efficiency and effectiveness of TRPL +K with other methods through experiments.
2. Background and related work. Since restarting techniques are the primary concern of this paper, we focus our background section on a common framework in which we can describe most current eigenmethods as well as the proposed one.
2.1. Thick restarting, locally optimal restarting, and preconditioning. The unrestarted Lanzcos method converges optimally in terms of the number of matrix-vector multiplications because it dynamically builds the optimal polynomial through an efficient three-term recurrence. In practice, rounding errors cause loss of orthogonality to previous Lanczos vectors, so we typically store all the Lanczos vectors and perform selective or partial reorthogonalization [30, 10]. Restarting is intended to reduce the storage requirements and computational cost of orthogonalization. After a maximum number of iteration vectors are stored, we compute the best desired approximations and restart. While limiting storage and computational costs per iteration, restarting inevitably impairs the optimality of unrestarted Lanczos since it discards part of the information. Various techniques [25, 36, 34, 17, 21] attempt to partially recover the lost information due to restarting.

Implicit restarting [33] performs this by dumping unwanted components (typically unwanted Ritz vectors) by applying the implicitly shifted QR. However, this technique is complicated to implement in a stable way [14, p. 136]. Thick restarting is mathematically equivalent to implicit restarting [37, 22] yet it is easier to implement in a stable way in the Lanczos [43], Arnoldi [22], and GD [37] methods. Thick restarting directly keeps the wanted Ritz vectors instead of dumping the unwanted ones from the basis. Moreover, with thick restarting it is straightforward to add arbitrary (nonKrylov) vectors to the restarted space. This is a key feature for our proposed method where we augment the restarted space with Ritz vectors from a previous cycle. Due to simplicity, numerically stability, and flexibility, thick restarting has been applied to various Krylov and GD (or JD) type methods for both eigenvalue and SVD problems [43, 22, 41, 20, 32, 37, 34, 1, 47].

The locally optimal restarting technique has been studied under different names in the literature such as locally optimal recurrence in LOBPCG [17], +K restarting in GD+K [36, 34], and Krylov subspace optimization [21]. There are different ways to justify the use of this technique. One is from an optimization viewpoint that extends the nonlinear conjugate gradient (CG) method for optimizing the Rayleigh quotient by a $R R$ procedure on the three successive iterates. Another viewpoint is that three
successive Lanczos iterates are sufficient to guarantee full orthogonality of the space. Yet another viewpoint relates the Lanczos iterates to the ones from a three-term recurrence of the linear CG on the JD correction equation. Regardless of the viewpoint, the idea of using Ritz vectors from both the current and the previous iteration has given rise to methods that converge near-optimally for seeking one eigenpair under limited memory, especially when combined with block methods or thick restarting [17, 34]. Rigorous analysis, however, has been difficult. In [36] the last viewpoint was analyzed providing some bounds on how well the locally optimal restarting matches the effects of a global optimization over the unrestarted space. In this paper, we provide an optimization viewpoint analysis that establishes the asymptotic global quasi optimality of our new method by quantifying the relative difference between the locally and a globally optimal Rayleigh quotients.

Preconditioning (inexact shift-invert) is crucial to improve the convergence of all these methods. For large problems, as exact matrix factorizations are prohibitive or infeasible, we focus on preconditioners such as incomplete ILU or $L D L^{T}$ factorization [11]. Although GD (or JD) type methods use a preconditioner to build a general, non-Krylov subspace, a few methods have been proposed to exploit a preconditioned Krylov subspace [11, 24, 41]. This paper further explores this line of research.
2.2. Comparison of subspaces of various methods. We use the following notation. A cycle refers to all the work a method performs between restarts and is denoted by a subscript. When present, a second subscript refers to the eigenvector index, e.g., $x_{i, 2}$ is the approximate eigenvector for the second smallest eigenvalue at the end of cycle $i$. A matrix or block of vectors followed by parenthesis uses MATLAB index notation. At restart, thick restarting keeps at least the $p$ wanted Ritz vectors. A method then continues building a basis $U$ with $m$ new vectors, which differentiates most methods. In +K restarting, $U$ is then augmented by $l$ Ritz vectors from the previous step (which previous step depends on the method). Thus, at the end of the $i$ th outer cycle, the basis is $U \in \mathbb{R}^{n \times q}$, where $q=p+m+l$ is the maximum basis size. For a given shift $\rho$, let $A_{\rho}=(A-\rho I), \widehat{A}_{\rho}=L^{-1}(A-\rho I) L^{-T}$, and a preconditioner $M=L^{-T} L^{-1} \approx(A-\rho I)^{-1}$. The shift is usually the Rayleigh quotient for some approximate eigenvector $x, \rho(x)=\frac{x^{T} A x}{x^{T} x}$. We use the RR procedure to extract approximate eigenpairs from $\operatorname{span}(U) . \mathcal{K}_{m}\left(A, u_{1}\right)$ refers to the Krylov subspace of dimension $m$ of $A$ with initial vector $u_{1}$.

TRLan [43] is mathematically equivalent to implicitly restarted Lanczos. At the end of the $i+1$ cycle, it forms a space $K_{T R L}$ which includes the $p$ wanted Ritz vectors obtained by RR at the end of the $i$ th cycle and $m$ additional Lanczos vectors starting from the last Lanczos vector $r=u_{i, p+m+1}$ of the $i$ th cycle.

$$
\begin{equation*}
K_{T R L}=\operatorname{span}\{\underbrace{x_{i, 1}, x_{i, 2}, \ldots, x_{i, p}}_{\text {Wanted Ritz vectors }}, \underbrace{r, A r, A^{2} r, \ldots, A^{m-1} r}_{\text {Lanczos iterations }}\} . \tag{2.1}
\end{equation*}
$$

$K_{T R L}$ always remains a Krylov space, but it allows for a more efficient implementation than the implicitly restarted Lanczos. TRLan is the only method we consider that cannot use preconditioning directly, but only through shift-invert.

LOBPCG [17] forms a subspace $K_{L}$ at the end of cycle $i+1$ from which it will use RR to compute $p$ approximate eigenpairs. The subspace is built by the following locally optimal recurrence with $m=p=l$ :

$$
\begin{equation*}
K_{L}=\operatorname{span}\{\underbrace{x_{i, 1}, x_{i, 2}, \ldots, x_{i, p},}_{\text {Wanted Ritz vectors }} \underbrace{M r_{1}, M r_{2}, \ldots, M r_{p},}_{\text {Preconditioned residual vectors }} \underbrace{x_{i-1,1}, \ldots, x_{i-1, p}}_{\text {Previous Ritz vectors }}\}, \tag{2.2}
\end{equation*}
$$

where $x_{i, j}$ is the $j$ th Ritz vector from cycle $i,\left\{r_{j}\right\}$ is its residual, and $x_{i-1, j}$ is the corresponding Ritz vector from cycle $i-1 .{ }^{1}$ Note that $\left\{M r_{j}\right\}$ are typically computed in a block form, not one at a time through an inner iteration.

GD $+\mathrm{K}[36,34]$ is an extension of GD where, at restart, $l$ number of previous Ritz vectors are used to augment the thick restarted basis. Thus at the end of cycle $i+1$ the subspace of size $q$ is

$$
\begin{equation*}
K_{G}=\operatorname{span}\{\underbrace{x_{i, 1}, \ldots, x_{i, p},}_{\text {Wanted Ritz vectors }} \underbrace{M r_{1}^{(p)}, \ldots, M r_{1}^{(p+m-1)}}_{\text {Preconditioned residual vectors }} \underbrace{x_{i, 1}^{(q-1)}, \ldots, x_{i, l}^{(q-1)}}_{\text {Previous Ritz vectors }}\} . \tag{2.3}
\end{equation*}
$$

Here, $x_{i, j}$ and $x_{i, j}^{(q-1)}$ are the Ritz vectors computed at cycle $i$ from the spaces $K_{G}$ with $\operatorname{size}\left(K_{G}\right)=q$ and $\operatorname{size}\left(K_{G}\right)=q-1$, respectively. $r^{(j)}$ denotes the residual vector of the targeted Ritz vector $x_{i+1,1}^{(j)}$ at inner iteration $j-p+1$ of the current cycle $i+1$, i.e., when $\operatorname{size}\left(K_{G}\right)=j$. Without preconditioning $(M=I)$ and with $l=0, \mathrm{GD}+\mathrm{K}$ is equivalent to TRLan. When $m=p=l=1, \mathrm{GD}+\mathrm{K}$ is equivalent to the locally optimal preconditioned conjugate gradient (LOPCG, or LOBPCG with $p=1$ ). A block version of $\mathrm{GD}+\mathrm{K}$ is also possible.

PL [24] employs a preconditioned Krylov inner iteration on $\widehat{A}_{\rho}=L^{-1}(A-\rho I) L^{-T}$ to build a basis $G$ of $K_{P}$ in (2.4), applies the RR of $\widehat{A}_{\rho}$ onto $G$ to find a primitive Ritz pair $(\theta, y)$, and converts the Ritz pair back $\left(\rho+\theta, L^{-T} G y\right)$ for the original eigenvalue problem. Additional eigenpairs are found one at a time. The size $m$ of the Krylov space varies dynamically.

$$
\begin{equation*}
K_{P}=\operatorname{span}\{\underbrace{L^{T} x_{i, 1},}_{\text {Wanted Ritz vector }} \underbrace{\widehat{A}_{\rho}\left(L^{T} x_{i, 1}\right), \widehat{A}_{\rho}^{2}\left(L^{T} x_{i, 1}\right), \ldots, \widehat{A}_{\rho}^{m}\left(L^{T} x_{i, 1}\right)}_{\text {Lanczos iterations }}\} . \tag{2.4}
\end{equation*}
$$

Without preconditioning PL is an explicitly restarted Lanczos with one vector (i.e., no thick restarting).

EigIFP [11] produces the same approximations as PL in exact arithmetic, but the application of the preconditioner does not have to be in factorized form. The basis $V$ of the search space $K_{F}$ built by EigIFP is related to the PL's basis as $U=L^{-T} G$. RR is performed by projecting $A_{\rho}=A-\rho I$ onto $U$, yielding the approximate eigenpair $(\rho+\theta, U y)$. Since all vectors are stored, $m=q-1$. Otherwise, EigIFP has the same limitations as PL.

$$
\begin{equation*}
K_{F}=\operatorname{span}\{\underbrace{x_{i, 1},}_{\text {Wanted Ritz vector }} \underbrace{M A_{\rho} x_{i, 1},\left(M A_{\rho}\right)^{2} x_{i, 1}, \ldots,\left(M A_{\rho}\right)^{m} x_{i, 1}}_{\text {Lanczos iterations }}\} . \tag{2.5}
\end{equation*}
$$

The GD+K method typically demonstrates faster convergence than the rest of the methods in terms of number of matrix-vector products because it combines both thick and locally optimal restarting and uses subspace acceleration to obtain the "best" Ritz vector at every inner iteration to improve by preconditioning. However, the faster convergence of $\mathrm{GD}+\mathrm{K}$ is at the cost of applying a more frequent RR procedure, which could be a quite expensive operation when the subspace is large [40]. Therefore, it is unclear if the GD+K method is still the method of choice in terms of the runtime when the matrix-vector operation is inexpensive. For large numbers of well-separated eigenvalues, however, a block method such as LOBPCG could also be competitive. On the other hand, PL and EigIFP can generate the inner Krylov space without the overhead of multiple RR projections required by GD+K.

[^1]3. Thick-restart preconditioned Lanczos $+\mathbf{K}$ method. The motivation of our proposed TRPL +K method is to extend the computationally efficient EigIFP method with the thick and locally optimal restarting techniques of GD+K. It can also be viewed as an extension of TRLan to allow for locally optimal restarting and preconditioning. We note that the JD method with CG as inner iteration and GD+K as the outer method is even more computationally efficient per step because the inner Krylov space does not need to be stored. However, the result of CG is a correction to a single Ritz vector that does not benefit convergence to nearby eigenpairs, which is left to the subspace acceleration of the outer iteration. With TRPL+K we hope that the inner iteration generates useful correction information to all $p$ required eigenvectors. For simplicity, we describe the method in detail for the standard eigenvalue problem only.

Extending EigIFP to include thick restarting is rather straightforward. Without preconditioning, the $K_{F}$ space in (2.5) is a Krylov space so it can be restarted as in TRLan (2.1). With preconditioning, $K_{F}$ is still a Krylov space but of the preconditioned matrix $M(A-\rho I)$. Note that if $\rho$ differs between cycles so do the matrices of the corresponding Krylov subspaces. At the end of cycle $i+1$, the RR must be performed on matrix $A$ or $A_{\rho}=A-\rho I$ to compute the new Ritz vectors $X_{i+1}=\left[x_{i+1,1}, \ldots, x_{i+1, p}\right]$. This implies that the thick restart vectors $X_{i}$ used in the basis of cycle $i+1$ do not form vectors of a Krylov sequence. After restart, we have $T=X_{i}^{T} A X_{i}=\operatorname{diag}\left(\left[\theta_{i, 1}, \ldots, \theta_{i, p}\right]\right)$, but we cannot use the TRLan relations for the subsequent Krylov vectors.

To efficiently build such an augmented Krylov space we can use a technique based on an FGMRES-like method [3] or on the GCRO method [7, 8]. We have followed a GCRO-like method to build a Krylov basis $G_{i+1}$ orthogonal to $X_{i}=\left[x_{i, 1}, \ldots, x_{i, p}\right]$, i.e., we build $\operatorname{span}\left(G_{i+1}\right)=\mathcal{K}_{m}\left(\left(I-X_{i} X_{i}^{T}\right) M(A-\rho I),\left(I-X_{i} X_{i}^{T}\right) M r_{1}\right)$, with $r_{1}=$ $(A-\rho I) x_{i, 1}$. This allows us to build the projection matrix $T=\left[X_{i}, G_{i+1}\right]^{T} A\left[X_{i}, G_{i+1}\right]$ without additional matrix-vector products. At step $j$ of the inner Krylov method, before we compute $G_{i+1}(:, j+1)$ we compute

$$
\begin{align*}
z & =A G_{i+1}(:, j) \\
T(1: p, p+j) & =X_{i}^{T} z  \tag{3.1}\\
T(p+1: p+j, p+j) & =G_{i+1}(:, 1: j)^{T} z
\end{align*}
$$

Then, we continue with the inner method $G_{i+1}(:, j+1)=\left(I-X_{i} X_{i}^{T}\right) M\left(z-\rho x_{i, 1}\right)$.
Employing locally optimal restarting to the above thick restarted preconditioned Lanczos is more involved. At cycle $i+1$, after the inner method concludes its $m$ steps, we perform RR to obtain $X_{i+1}$ using the space $U_{i+1}(:, 1: m+p)=\left[X_{i}, G_{i+1}\right]$. We want to augment this space with some "previous" directions, $X^{p r e v}$. A choice similar to $\mathrm{GD}+\mathrm{K}$ does not work. $\mathrm{GD}+\mathrm{K}$ uses as $X^{\text {prev }}$ the Ritz vectors from the penultimate step of the Krylov method before restart, i.e., the Ritz vectors from the subspace $U_{i}(:, 1: p+m-1)=\left[X_{i-1}, G_{i}(:, 1: m-1)\right]$. The idea is that the optimal projection over $U_{i}(:, 1: p+m+1)$ (i.e., the next step of the unrestarted method) can be approximated through the Ritz vectors of the last two iterations and the residual. But our method does not optimize over the $X^{\text {prev }}$ directions to expand the basis at cycle $i+1$; it builds a Krylov space.

If we assume that our inner method was a polynomial returning only one vector $s(A) r_{1}$ (not the entire $G_{i+1}$ space of size $m$ ) to be used in the outer RR and that we were seeking $p=1$ eigenpair, the outer method would be similar to LOBPCG on the operator $s(A)$. For this operator, the choice for locally optimal restarting would
be $X^{\text {prev }}=X_{i-1}$, i.e., the Ritz vector at the beginning of the previous cycle. We take the liberty to use LOBPCG's choice $X^{\text {prev }}$ for our method, which we now call TRPL +K . The space at the end of cycle $i+1$ should include $X_{i}, G_{i+1}$, and $X_{i-1}$.

The order of the three blocks in the search subspace deserves a careful discussion. Since $X_{i}$ is produced from a basis that includes $X_{i-1}$ at cycle $i$, it is possible to orthogonalize the two sets implicitly and compute their interaction on the projection matrix $T_{i+1}$ without matrix-vector products. Then, by augmenting our space not only with $X_{i}$ but with $\operatorname{span}\left(\left[X_{i}, X_{i-1}\right]\right)$, we could build $G_{i+1}$ as described above. However, we noticed experimentally that this choice does not perform well.

We observed much better convergence if the previous vectors $X_{i-1}$ were included in the basis after $G_{i+1}$ was computed. The disadvantage is that $X_{i-1}$ has to be explicitly orthogonalized against the rest of the basis vectors and $l$ matrix-vector products have to be performed to compute the resulting $T$. However, the relative expense is small for large $m$ (inner iterations), and typically $l=1$ gives very close to optimal convergence while more previous vectors give little additional improvement.

We can now describe the space that TRPL +K builds. Let $X_{i}=\left[x_{i, 1}, \ldots, x_{i, p}\right]$, $C=\left(I-X_{i} X_{i}^{T}\right) M\left(A-\rho_{i} I\right)$, and $u_{1}=C x_{i, 1}$. Then, at the end of cycle $i+1$ TRPL +K computes the Ritz vectors from the subspace

$$
\begin{equation*}
K_{T R P L+K}=\operatorname{span}\{\underbrace{x_{i, 1}, \ldots, x_{i, p}}_{\text {Wanted Ritz vector }} \underbrace{u_{1}, C u_{1}, \ldots, C^{m-1} u_{1}}_{\text {Lanczos iterations }} \underbrace{x_{i-1,1}, \ldots, x_{i-1, l}}_{\text {Previous Ritz vector }}\} \tag{3.2}
\end{equation*}
$$

Algorithm 1 summarizes TRPL +K for finding $p$ smallest eigenpairs of $A$. For simplicity we also let the minimum thick restart size be $p$. However, the algorithm

```
Algorithm 1. Thick-restart preconditioned Lanczos +K method.
    Input: matrix \(A \in \mathbb{R}^{n \times n}\), preconditioner \(M \in \mathbb{R}^{n \times n}, p\) the number of desired
    eigenpairs, \(X=\left[x_{1} \ldots x_{p}\right]\) any available initial approximations to the \(p\) eigenpairs,
    \(x_{i}^{T} x_{j}=\delta_{i j}\), maximum basis size \(q\), maximum number of retained previous vectors
    \(l\), maximum number of cycles maxIter
    Output: approximate eigenvalues \(\theta_{1}, \ldots, \theta_{p}\), and eigenvectors \(x_{1}, \ldots, x_{p}\).
    Perform RR on the orthonormal set \(U=X=\left[x_{1}, \ldots, x_{p}\right], T(1: p, 1: p) \leftarrow X^{T} A X\)
    Let \(Y \Theta Y^{T}=X^{T} A X\) be the eigendecomposition of \(X^{T} A X\). Reset \(U, X \leftarrow U Y\)
    (the \(p\) Ritz vectors) and \(T \leftarrow \Theta\). Set target \(t=1, \rho \leftarrow x_{t}^{T} A x_{t}, X^{\text {prev }} \leftarrow[]\)
    for \(k=1:\) maxIter do
        Use Lanczos (3.1) to generate \(G\) containing the orthonormal basis vectors for
        \(\mathcal{K}_{m}\left(\left(I-X X^{T}\right) M(A-\rho I), u_{1}\right)\) and the projection \(T(1: p+m, p+1: p+m)\)
        Set \(U(:, 1: p+m) \leftarrow[X, G]\)
        if \(k>1\) then
            Orthogonalize \(X^{\text {prev }}\) against \(U(:, 1: p+m)\) to build \(U(:, p+m+1: q)\)
            Compute the remaining \(T\) such that \(T=U^{T} A U\)
        end if
        \(X^{\text {prev }}=\left[x_{1}^{\text {prev }}, \ldots, x_{l}^{\text {prev }}\right] \leftarrow\left[x_{t}, \ldots, x_{t+l-1}\right] \quad(\) up to a \(\min (t+l-1, p))\)
        Compute eigenpairs of \(T, T=Y \Theta Y^{T}\) and Ritz pairs: \(\left(\theta_{i}, x_{i}=U Y(:, i)\right)\)
        if \(\left\|A x_{t}-\theta_{t} x_{t}\right\| \leq \epsilon\) then
            flag \(\left(\theta_{t}, x_{t}\right)\), remove \(x_{1}^{\text {prev }}\) from \(X^{\text {prev }}\), and advance target \(t=t+1\)
        end if
        Set \(\rho \leftarrow \theta_{t}, U(:, 1: p+1) \leftarrow\left[U Y(:, 1: p), M\left(A x_{t}-\rho x_{t}\right)\right], X=U(:, 1: p)\)
    end for
```

can thick restart with any size $\hat{p}>p$. Note also that the preconditioner $M$ may change at every cycle with the goal to approximate $\left(A-\rho_{i} I\right)^{-1}$.

Finally, we note that the algorithm can be easily extended to the generalized eigenvalue problem, $A v=\lambda B v$ with $B$ symmetric positive definite. We forgo this description for simplicity. The analysis in the next section is based on a simplified version of our algorithm for the generalized eigenvalue problem, aiming to compute only the lowest eigenpair $\left(\lambda_{1}, v_{1}\right)(p=1)$, replacing the previous Ritz vector $x_{i-1,1}$ in (3.2) with a search direction connecting $x_{i-1,1}$ and $x_{i, 1}$ in two consecutive cycles. We show the algorithm with soft locking but hard locking can also be applied.
4. Asymptotic convergence analysis of a $\mathrm{PL}+1$ method. In this section, we consider a PL +1 method (i.e., TRPL +K with $p=1=l$ ) for computing the lowest eigenpair $\left(\lambda_{1}, v_{1}\right)$ of $(A, B)$, establishing an asymptotic global quasi optimality of the method.
4.1. Preliminaries. Consider the matrix pencil $(A, B)$, where $A, B \in \mathbb{R}^{n \times n}$ are symmetric, typically large and sparse, and $B$ is positive definite. Let $\lambda_{1}<\lambda_{2} \leq$ $\ldots \leq \lambda_{n}$ and $\left\{v_{i}\right\}$ be the eigenvalues and eigenvectors of the matrix pencil, such that $A v_{i}=\lambda_{i} B v_{i},\left(v_{i}, v_{j}\right)_{B}=v_{i}^{T} B v_{j}=\delta_{i j},\left\|v_{i}\right\|_{B}=1,1 \leq i, j \leq n$.

Let $x$ be an approximation to $v_{1}$, the eigenvector associated with the lowest eigenvalue $\lambda_{1}$, with the decomposition

$$
\begin{equation*}
x=v_{1} \cos \theta+f \sin \theta, \text { where } \theta \neq 0, f \perp B v_{1}, \text { and }\|f\|_{B}=1 \tag{4.1}
\end{equation*}
$$

This suggests that $\|x\|_{B}=\left(x^{T} B x\right)^{\frac{1}{2}}=\left(\left\|v_{1}\right\|_{B}^{2} \cos ^{2} \theta+\|f\|_{B}^{2} \sin ^{2} \theta\right)^{\frac{1}{2}}=1$. Since $f \perp B v_{1}$, it has the form $f=\sum_{j=2}^{n} s_{j} v_{j}$, where the scalars $\left\{s_{j}\right\}_{j=2}^{n}$ satisfy $\sum_{j=2}^{n} s_{j}^{2}=$ $\sum_{j=2}^{n} s_{j}^{2}\left\|v_{j}\right\|_{B}^{2}=\|f\|_{B}^{2}=1$. The Rayleigh quotient of $x$ is hence

$$
\rho(x)=\frac{x^{T} A x}{x^{T} B x}=x^{T} A x=v_{1}^{T} A v_{1} \cos ^{2} \theta+f^{T} A f \sin ^{2} \theta=\lambda_{1} \cos ^{2} \theta+\rho(f) \sin ^{2} \theta,
$$

where $\rho(f)=\frac{f^{T} A f}{f^{T} B f}=f^{T} A f=\sum_{j=2}^{n} s_{j}^{2} \lambda_{j} \in\left[\lambda_{2}, \lambda_{n}\right]$.
Proposition 4.1. Let $x$ be a vector with $\|x\|_{B}=1$. The gradient and the Hessian of $\frac{1}{2} \rho(x)$ with respect to $x$, respectively, are
(4.2) $\nabla \frac{1}{2} \rho(x)=\frac{1}{x^{T} B x}(A-\rho(x) B) x=A x-\rho(x) B x \quad$ and

$$
\begin{equation*}
\nabla^{2} \frac{1}{2} \rho(x)=\frac{1}{x^{T} B x}\left\{A-\rho(x) B-\frac{2}{x^{T} B x}(A x-\rho(x) B x)(B x)^{T}-\frac{2}{x^{T} B x} B x(A x-\rho(x) B x)^{T}\right\} \tag{4.3}
\end{equation*}
$$

Proof. The proof is done by letting $T(\rho)=\rho B-A$ in [38, Proposition 3.1].
For small $\theta, \rho(x)$ is a second order approximation to $\lambda_{1}$, i.e.,

$$
\begin{equation*}
\rho(x)-\lambda_{1}=\sin ^{2} \theta\left(\rho(f)-\lambda_{1}\right) \tag{4.4}
\end{equation*}
$$

and hence $\nabla \frac{1}{2} \rho(x)$, i.e., the eigenresidual associated with $x$, is

$$
\begin{aligned}
& A x-\rho(x) B x=(A-\rho(x) B)\left(v_{1} \cos \theta+f \sin \theta\right) \\
= & \left(\lambda_{1}-\rho(x)\right) \cos \theta B v_{1}+\sin \theta(A-\rho(x) B) f \\
= & \sin \theta\left[\sum_{j=2}^{n} s_{j}\left(\lambda_{j}-\rho(x)\right) B v_{j}-\sin \theta \cos \theta\left(\rho(f)-\lambda_{1}\right) B v_{1}\right] .
\end{aligned}
$$

Since $\lambda_{1}<\lambda_{2}$ and $\sum_{j=2}^{n} s_{j}^{2}=1, A f-\rho(x) B f=\sum_{j=2}^{n} s_{j}\left(\lambda_{j}-\rho(x)\right) B v_{j}$ will not vanish as $\theta \rightarrow 0$ and $\rho(x) \rightarrow \lambda_{1}$. Therefore, for sufficiently small $\theta$,

$$
\begin{equation*}
(1-\delta) \sin \theta\|A f-\rho(x) B f\| \leq\|A x-\rho(x) B x\| \leq(1+\delta) \sin \theta\|A f-\rho(x) B f\| \tag{4.5}
\end{equation*}
$$

for some small $\delta>0$ independent of $\theta$, or simply $\|A x-\rho(x) B x\|=\mathcal{O}(\sin \theta)$.
4.2. A $\mathbf{P L}+1$ method. The framework of a $\mathrm{PL}+1$ method is summarized in Algorithm 2. This is a special case of Algorithm 1 for $p=l=1$ but extends it to the generalized eigenvalue problem. It can also be considered an extension of the PL method enhanced with the search direction adopted in LOPCG. It is important to note that when $X$ is a single vector, $\left[X_{k}, G_{k+1}\right]$ in (3.2) is a Krylov space and $\operatorname{span}\left\{x_{k, 1}, u_{1}, C u_{1}, \ldots, C^{m-1} u_{1}\right\}=\operatorname{span}\left\{x_{k, 1}, A x_{k, 1}, \ldots A^{m} x_{k, 1}\right\}$ (see Lemma 4.1 in [36]). Since we only look for one eigenvalue, in the rest of this section we drop the second subscript notation and use the subscript $k$ to represent the cycle, hence the simpler form in Algorithm 2.

```
Algorithm 2. PL +1 method for computing \(\left(\lambda_{1}, v_{1}\right)\) of \((A, B)\).
    Choose a symmetric positive definite preconditioner \(M \approx(A-\sigma B)^{-1}\left(\sigma<\lambda_{1}\right)\),
    tolerance \(\delta>0, m>0\)
    Choose vector \(x_{0}\) with \(\left\|x_{0}\right\|_{B}=1\), set \(\rho_{0}=\frac{x_{0}^{T} A x_{0}}{x_{0}^{T} B x_{0}}\) and \(r_{0}=\left(A-\rho_{0} B\right) x_{0}\).
    for \(k=0,1, \ldots\), until the convergence, i.e., \(\left\|r_{k}\right\| \leq \delta\) do
        Form \(G_{k}\) containing \(B\)-normalized basis vectors of \(M \mathcal{K}_{m}\left(\left(A-\rho_{k} B\right) M, r_{k}\right)\)
        \(=\operatorname{span}\left\{M\left(A-\rho_{k} B\right) x_{k},\left[M\left(A-\rho_{k} B\right)\right]^{2} x_{k}, \ldots,\left[M\left(A-\rho_{k} B\right)\right]^{m} x_{k}\right\}\)
        Form \(Q_{k}=\left\{\begin{array}{ll}{\left[x_{k}, G_{k}\right]} & (k=0) \\ {\left[x_{k}, G_{k}, p_{k-1}\right]} & (k>0)\end{array}\right.\), perform the RR projection and
        solve \(Q_{k}^{T} A Q_{k} w=\rho Q_{k}^{T} B Q_{k} w\) for the lowest primitive Ritz pair \(\left(\rho_{k}, w_{k}\right)\).
        \(g_{k}=G_{k} w_{k}(2: m+1)\)
        \(\widetilde{x}_{k+1}=Q_{k} w_{k}= \begin{cases}x_{k} w_{k}(1)+g_{k} & (\text { if } k=0), \\ x_{k} w_{k}(1)+g_{k}+p_{k-1} w_{k}(m+2) & (\text { if } k>0),\end{cases}\)
        \(x_{k+1}=\frac{\widetilde{x}_{k+1}}{\left\|\tilde{x}_{k+1}\right\|_{B}}, \rho_{k+1}=\frac{x_{k+1}^{T} A x_{k+1}}{x_{k+1}^{T} B x_{k+1}}\), and \(r_{k+1}=\left(A-\rho_{k+1} B\right) x_{k+1}\).
        \(\widetilde{p}_{k}=\left\{\begin{array}{ll}g_{k} & (k=0) \\ g_{k}-\frac{p_{k-1}^{T}\left(A-\rho_{k-1} B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho_{k-1} B\right) p_{k-1}} p_{k-1} & (k>0)\end{array}, p_{k}=\frac{\widetilde{p}_{k}}{\left\|\tilde{p}_{k}\right\|_{B}}\right.\).
    end for
```

In each outer iteration $k(k>0)$, at steps 4 and 5 , an augmented Krylov subspace $\operatorname{range}\left(Q_{k}\right)=\operatorname{span}\left\{x_{k}\right\}+M \mathcal{K}_{m}\left(\left(A-\rho_{k} B\right) M, r_{k}\right)+\operatorname{span}\left\{p_{k-1}\right\}$, i.e.,

$$
\begin{equation*}
\mathcal{K}_{m+1}\left(M\left(A-\rho_{k} B\right), x_{k}\right)+\operatorname{span}\left\{p_{k-1}\right\}, \tag{4.6}
\end{equation*}
$$

is formed as the subspace for the RR projection. Algorithm 2 is slightly different from the variant used in practice at step 7 , where a commonly adopted approach sets $p_{k}=g_{k}+p_{k-1} w_{k}(m+2)$, using the $(m+2)$ nd element of primitive Ritz vector $w_{k}$ as the coefficient for $p_{k-1}$. Our choice of such a particular linear combination makes it easy to show the near conjugacy between $p_{k-1}$ and $p_{k}$.

We present a preliminary convergence result for Algorithm 2, which is essentially a restatement of [11, Theorem 3.4]. Here we incorporate preconditioning and note the fact that the space for projection used in [11], $\mathcal{K}_{m+1}\left(M\left(A-\rho_{k} B\right), x_{k}\right)$, is a subspace of the one constructed by Algorithm 2 in (4.6).

ThEOREM 1. Let $\lambda_{1}<\lambda_{2} \leq \cdots \leq \lambda_{n}$ be the eigenvalues of $(A, B),\left(\rho_{k}, x_{k}\right)$ be the $k$ th iterate of Algorithm 2, and $L D L^{T} \approx\left(A-\rho_{k} B\right)$ be a preconditioner where $D$ is a diagonal matrix of $\pm 1$ elements. Assume that $C_{k}=L^{-1}\left(A-\rho_{k} B\right) L^{-T}$ has eigenvalues $\sigma_{1}<\sigma_{2} \leq \cdots \leq \sigma_{n}$ and satisfies $C_{k} u_{1}=\sigma_{1} u_{1}$ with $\left\|u_{1}\right\|_{2}=1$. If $\lambda_{1}<\rho_{k}<\lambda_{2}$, then $\rho_{k+1}-\lambda_{1} \leq\left(\rho_{k}-\lambda_{1}\right) \epsilon_{m}^{2}+\left(\rho_{k}-\lambda_{1}\right)^{\frac{3}{2}} \epsilon_{m}\left(\frac{\left\|L^{-1} B L^{-T}\right\|}{\sigma_{2}}\right)^{\frac{1}{2}}+\delta_{k}$, where $0<\delta_{k} \equiv$ $\rho_{k}-\lambda_{1}+\frac{\sigma_{1}}{u_{1}^{T} L^{-1} B L^{-T} u_{1}}=\mathcal{O}\left(\left(\rho_{k}-\lambda_{1}\right)^{2}\right)$, and $\epsilon_{m}=\min _{p \in \mathcal{P}_{m}, p\left(\sigma_{1}\right)=1} \max _{2 \leq i \leq n}\left|p\left(\sigma_{i}\right)\right|$, with $\mathcal{P}_{m}$ denoting the set of all polynomials of degree not greater than $m$.

Remark. The value of $\epsilon_{m}$ depends on the quality of the preconditioner $M$ approximating $A-\rho_{k} B$. If $M=A-\sigma B$, then $\left(A-\rho_{k} B\right) V=B V\left(\Lambda-\rho_{k} I\right)$ and $M V=B V(\Lambda-\sigma I)$, where $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{n}\right)$ is the diagonal matrix of the eigenvalues of $(A, B)$ and $V=\left[v_{1}, \ldots, v_{n}\right]$ contains the corresponding eigenvectors. It follows that $M\left(A-\rho_{k} B\right) V=V(\Lambda-\sigma I)^{-1}\left(\Lambda-\rho_{k} I\right)$. That is, $v_{i}$ is an eigenvector of $M\left(A-\rho_{k} B\right)$ associated with eigenvalue $\sigma_{i}=\frac{\lambda_{i}-\rho_{k}}{\lambda_{i}-\sigma}$. Since $C_{k}=L^{-1}\left(A-\rho_{k} B\right) L^{-T}$ and $L^{-T} L^{-1}\left(A-\rho_{k} B\right)$ have identical spectrum, $\sigma_{i}$ is also an eigenvalue of $C_{k}$. Therefore, if $\sigma$ and $\rho_{k}$ are close to $\lambda_{1}$, then $\sigma_{2}, \ldots, \sigma_{n}$ are all close to 1 , and hence $\epsilon_{m}=\min _{p \in \mathcal{P}_{m}, p\left(\sigma_{1}\right)=1} \max _{2 \leq i \leq n}\left|p\left(\sigma_{i}\right)\right|$ would be fairly small with a small value of $m$, indicating a fast rate of convergence of the outer iteration.
4.3. Global quasi optimality. Theorem 1 shows that Algorithm 2 converges at least linearly with an asymptotic factor no greater than $\epsilon_{m}^{2}$. Our goal is to explore the role of the search directions $p_{k}$ (the " +1 " strategy), which helps the algorithm achieve a greatly improved convergence rate associated with global quasi optimality. The global quasi optimality is defined as follows.

Definition 4.2. Consider an iterative method for computing the lowest eigenpair $\left(\lambda_{1}, v_{1}\right)$ of a real symmetric matrix pencil $(A, B)$ with positive definite $B$. Let $x_{0}$ be the starting vector, $x_{k}$ be the approximation obtained at step $k$, and $\theta_{k}=\angle\left(x_{k}, v_{1}\right)_{B}=$ $\cos ^{-1} \frac{\left(v_{1}, x_{k}\right)_{B}}{\left\|v_{1}\right\|_{B}\left\|x_{k}\right\|_{B}}$ be the error angle of $x_{k}$. Let $U_{1}, U_{2}, \ldots$ be a sequence of subspaces of increasing dimension, such that for each $k \geq 1, x_{k} \in U_{k}$, and $U_{i} \subset U_{j}$ for all $1 \leq i<j$. Let $y_{k}^{*} \in U_{k}$ be the global minimizer of the Rayleigh quotient $\rho(x)=\frac{x^{T} A x}{x^{T} B x}$ in $U_{k}$. Then the iterate $x_{k}$ achieves global quasi optimality if

$$
\begin{equation*}
\lim _{\theta_{0} \rightarrow 0} \frac{\rho\left(x_{k}\right)-\rho\left(y_{k}^{*}\right)}{\rho\left(x_{k}\right)-\lambda_{1}}=0 \tag{4.7}
\end{equation*}
$$

4.3.1. Linear convergence assumption. To show the global quasi optimality of Algorithm 2, we first make an assumption of its precisely linear convergence.

Assumption 4.3. Assume that Algorithm 2 starting with initial $\rho\left(x_{0}\right) \in\left(\lambda_{1}, \lambda_{2}\right)$ converges precisely linearly (not superlinearly or faster) to $\lambda_{1}$; in other words, there exist constants $0<\underline{\xi}<\bar{\xi}<1$, independent of the progress of Algorithm 2, such that

$$
\begin{equation*}
\underline{\xi}^{k-j}\left(\rho_{j}-\lambda_{1}\right) \leq \rho_{k}-\lambda_{1} \leq \bar{\xi}^{k-j}\left(\rho_{j}-\lambda_{1}\right) \text { for all } 0 \leq j<k \tag{4.8}
\end{equation*}
$$

The assumption on the existence of a lower bound $\underline{\xi}$ on the convergence factor is realistic. Given $M \approx(A-\sigma B)^{-1}$ with a fixed $\sigma<\lambda_{1}$, and a fixed dimension $m$ of the preconditioned Krylov subspace, extensive experiments suggest that Algorithm 2 exhibits simply linear convergence as the outer iteration proceeds.

Assumption 4.3 has an equivalent form in terms of angles. Consider two iterates $x_{j}=v_{1} \cos \theta_{j}+f_{j} \sin \theta_{j}$ and $x_{k}=v_{1} \cos \theta_{k}+f_{k} \sin \theta_{k}, 0 \leq j<k$, where $f_{j}, f_{k} \perp B v_{1}$,
$\left\|f_{j}\right\|_{B}=\left\|f_{k}\right\|_{B}=1$, and $\left\|x_{j}\right\|_{B}=\left\|x_{k}\right\|_{B}=1$. From (4.4), we have (4.8) equivalent to $\underline{\xi}^{k-j} \sin ^{2} \theta_{j}\left(\rho\left(f_{j}\right)-\lambda_{1}\right) \leq \sin ^{2} \theta_{k}\left(\rho\left(f_{k}\right)-\lambda_{1}\right) \leq \bar{\xi}^{k-j} \sin ^{2} \theta_{j}\left(\rho\left(f_{j}\right)-\lambda_{1}\right)$, i.e.,

$$
\begin{equation*}
\sqrt{\xi}^{k-j} \sqrt{\frac{\rho\left(f_{j}\right)-\lambda_{1}}{\rho\left(f_{k}\right)-\lambda_{1}}} \sin \theta_{j} \leq \sin \theta_{k} \leq \sqrt{\xi}^{k-j} \sqrt{\frac{\rho\left(f_{j}\right)-\lambda_{1}}{\rho\left(f_{k}\right)-\lambda_{1}}} \sin \theta_{j} . \tag{4.9}
\end{equation*}
$$

4.3.2. Relevant spaces. To study the global quasi optimality, we define

$$
\begin{align*}
W_{k}= & \operatorname{span}\left\{g_{0}, \ldots, g_{k-1}\right\}=\operatorname{span}\left\{M d_{0}, \ldots, M d_{k-1}\right\}, \text { where }  \tag{4.10}\\
& d_{\ell} \in \mathcal{K}_{m}\left(\left(A-\rho_{\ell} B\right) M, r_{\ell}\right) \text {, and } \\
U_{k}= & \operatorname{span}\left\{x_{0}\right\}+W_{k} . \tag{4.11}
\end{align*}
$$

Lemma 2. For all $k \geq 1, W_{k}=\operatorname{span}\left\{p_{0}, \ldots, p_{k-1}\right\}$ and $x_{k} \in U_{k}$.
The proof is straightforward by induction and omitted. Next, we make an important assumption about $W_{k}$ for the subsequent analysis.

Assumption 4.4. Assume that there is a constant $\delta>0$, independent of $\theta_{0}=$ $\angle\left(x_{0}, v_{1}\right)_{B}$, such that $\angle\left(v_{1}, W_{k}\right) \geq \delta$ for all $k \geq 1$.

Remark. The above assumption is guaranteed to hold if the preconditioner $M$ is equipped with the projector $P=I-\frac{x_{0} x_{0}^{T} B}{x_{0}^{T} B x_{0}}$. That is, let $\mathbb{M}=P^{T} M P$ be the preconditioner for Algorithm 2 such that range $\left(\mathbb{M}^{\dagger}\right)=\left(\operatorname{span}\left\{B x_{0}\right\}\right)^{\perp}$. If $x_{0}$ is sufficiently close to $v_{1}$, then $W_{k} \subset \operatorname{range}\left(\mathbb{M}^{\dagger}\right) \approx \operatorname{span}\left\{v_{2}, \ldots, v_{n}\right\}$, and hence $\angle\left(v_{1}, W_{k}\right) \geq \delta$. The JD method uses this strategy to enhance the robustness of the convergence.
4.3.3. Preliminary results. We present a few preliminary results useful for the proof of the main theorems in the next section. For the sake of readability, we move most of the technical proof of our results to the appendix.

Lemma 3. Let $x \approx v_{1}$ in (4.1), and let $p$ be a descent direction for $\rho(x)$ such that $(p, \nabla \rho(x))<0$. Assume that there exists a $\delta>0$ independent of $\theta=\angle\left(x, v_{1}\right)_{B}$ such that $\rho(p)-\rho(x) \geq \delta$. Then the optimal step size $\alpha^{*}$ minimizing $\rho(x+\alpha p)$ is the unique or the smaller positive root of $a(x, p) \alpha^{2}+b(x, p) \alpha+c(x, p)=0$, where

$$
\begin{align*}
a(x, p) & =\left(p^{T} A p\right)\left(p^{T} B x\right)-\left(p^{T} B p\right)\left(p^{T} A x\right)=\|p\|_{B}^{2} x^{T}(A-\rho(p) B) p,  \tag{4.12}\\
b(x, p) & =\left(p^{T} A p\right)\left(x^{T} B x\right)-\left(p^{T} B p\right)\left(x^{T} A x\right)=\|x\|_{B}^{2} p^{T}(A-\rho(x) B) p \\
& =\|x\|_{B}^{2}\|p\|_{B}^{2}(\rho(p)-\rho(x)) \geq\|x\|_{B}^{2}\|p\|_{B}^{2} \delta>0, \text { and } \\
c(x, p) & =\left(p^{T} A x\right)\left(x^{T} B x\right)-\left(p^{T} B x\right)\left(x^{T} A x\right)=\|x\|_{B}^{2} p^{T}(A-\rho(x) B) x<0 .
\end{align*}
$$

Lemma 4. Let $x$ be an approximation to $v_{1}$ with decomposition (4.1), $p$ be a descent direction for $\rho(x)$, and $\delta>0$ be a constant independent of $\theta=\angle\left(x, v_{1}\right)_{B}$, such that $\rho(p)-\rho(x) \geq \delta>0$, and $\alpha^{*}$ the optimal step size minimizing $\rho(x+\alpha p)$. Then for sufficiently small $\theta$, with $a, b$, and $c$ defined in (4.12),

$$
\frac{-\frac{c^{2}}{b b}-\frac{a c^{3}}{3 b^{3}}+\mathcal{O}\left(c^{4}\right)}{\min \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \leq \rho\left(x+\alpha^{*} p\right)-\rho(x) \leq \frac{-\frac{c^{2}}{2 b}-\frac{a c^{3}}{3 b^{3}}+\mathcal{O}\left(c^{4}\right)}{\max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}},
$$

and hence $\rho\left(x+\alpha^{*} p\right)-\rho(x)=-\mathcal{O}\left(\sin ^{2} \theta\right) \mathcal{O}\left(\cos ^{2} \angle(p, \nabla \rho(x))\right)$.
Lemma 5. Let $j \geq 0$ be fixed and $k>j$ be a variable integer. Consider two iterates $x_{j}=v_{1} \cos \theta_{j}+f_{j} \sin \theta_{j} \in U_{j}, x_{k}=v_{1} \cos \theta_{k}+f_{k} \sin \theta_{k} \in U_{k}$ computed by

Algorithm 2, where $f_{j}, f_{k} \perp B v_{1},\left\|f_{j}\right\|_{B}=\left\|f_{k}\right\|_{B}=1,\left\|x_{j}\right\|_{B}=\left\|x_{k}\right\|_{B}=1$, with sufficiently small $\theta_{j}, \theta_{k}$ such that $\rho\left(x_{k}\right) \leq \rho\left(x_{j}\right)<\lambda_{2}$. Consider a decomposition

$$
\begin{equation*}
x_{k}=\beta_{j k} x_{j}+\gamma_{j k} g_{j k} \tag{4.13}
\end{equation*}
$$

with $\beta_{j k}, \gamma_{j k}>0,\left\|g_{j k}\right\|_{B}=1, g_{j k} \in \operatorname{span}\left\{p_{j-1}, g_{j}, \ldots, g_{k-1}\right\}(j>0)$, or $g_{0 k} \in W_{k}=$ $\operatorname{span}\left\{g_{0}, \ldots, g_{k-1}\right\}(j=0)$. Under Assumption 4.4, let $\mu_{j k}=g_{j k}^{T} B x_{j}$, and we have

$$
\begin{equation*}
\gamma_{j k}=\mathcal{O}\left(\sin \theta_{j}\right), \beta_{j k}=1-\mu_{j k} \gamma_{j k}-\frac{1}{2}\left(1-\mu_{j k}^{2}\right) \gamma_{j k}^{2}+\mathcal{O}\left(\gamma_{j k}^{4}\right)=1-\mathcal{O}\left(\sin \theta_{j}\right) \tag{4.14}
\end{equation*}
$$

4.3.4. Main theorems. We are ready to prove the global quasi optimality of Algorithm 2. To this end, we will show the following results step by step:

1. If the search directions $\left\{p_{0}, \ldots, p_{k}\right\}$ of Algorithm 2 are approximately conjugate, then $x_{k+1} \in U_{k+1}$ is sufficiently close to the global minimizer in $U_{k+1}$ as long as $x_{k} \in U_{k}$ is sufficiently close to the global minimizer in $U_{k}$.
2. Any two consecutive search directions $p_{k-1}$ and $p_{k}$ are approximately conjugate, and so are $p_{0}$ and $p_{2}$ (hence, $x_{3}$ is globally quasi-optimal).
3. If $x_{k}$ is globally quasi-optimal, then $r_{k}$ is nearly orthogonal to $W_{k}$; in fact, $\lim _{\theta_{0} \rightarrow 0} \cos \angle\left(r_{k}, W_{k}\right)=0$ (we can hence define $\left.\cos \angle\left(r_{k}, W_{k}\right)\right|_{\theta_{0}=0}=0$ ).
4. Assume that $\cos \angle\left(r_{k}, W_{k}\right)$ is differentiable at $\theta_{0}=0$. If $x_{k} \in U_{k}$ is globally quasi-optimal, then $p_{k}$ is approximately conjugate to $\left\{p_{0}, \ldots, p_{k-1}\right\}$, and hence $x_{k+1} \in U_{k+1}$ is also globally quasi-optimal, as a result of induction.
First, we show that a set of approximately conjugate search directions guarantee that the quality of the iterate of Algorithm 2 at (outer) step $k$ for approximating the corresponding global minimizer can be extended to step $k+1$ with a possible very small deterioration on the order of $\mathcal{O}\left(\sin \theta_{k+1}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right)$.

Theorem 6. Let $\left\{x_{k}\right\}$ be the iterates of Algorithm $2, r_{k}=\left(A-\rho\left(x_{k}\right) B\right) x_{k}$ the residual, and $\left\{p_{k}\right\}$ the $B$-normalized search directions. For a given $k \geq 0$, consider all vectors of the form $z=y+\alpha p_{k} \in U_{k+1}$, where $y=\beta_{y} x_{0}+\gamma_{y} g_{y} \in U_{k}$ with $\|y\|_{B}=\left\|g_{y}\right\|_{B}=1$, and $g_{y} \in W_{k}$, satisfying $\rho(z) \leq \rho(y) \leq \rho\left(x_{0}\right)$. Assume that $\left\{p_{k}\right\}$ are pairwise approximately conjugate, i.e., $p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{j}=\mathcal{O}\left(\sin \theta_{j}\right)(0 \leq j<k)$. Let $y^{*}$ and $z^{*}$ be the global minimizer of $\rho(\cdot)$ in $U_{k}$ and $U_{k+1}$, respectively. Then

$$
\rho\left(x_{k+1}\right)-\rho\left(z^{*}\right) \leq \rho\left(x_{k}\right)-\rho\left(y^{*}\right)+\mathcal{O}\left(\sin \theta_{k+1}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right) .
$$

To prove the global quasi optimality of Algorithm 2, it is hence crucial to show that the $B$-normalized search directions are pairwise approximately conjugate, i.e.,

$$
p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{j}=\mathcal{O}\left(\sin \theta_{j}\right)
$$

for all integers $0 \leq j<k$. To achieve this, our second step is to show that any two consecutive search directions $p_{k-1}$ and $p_{k}$ are approximately conjugate, and so are $p_{0}$ and $p_{2}$. We will establish the complete near conjugacy in Theorem 10.

Lemma 7. The $B$-normalized search directions of Algorithm 2 satisfy

$$
p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{k-1}=\mathcal{O}\left(\sin ^{2} \theta_{k-1}\right) \quad \text { for all } k \geq 1
$$

To show the complete near conjugacy $p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{j}=\mathcal{O}\left(\sin \theta_{j}\right)$ for all $0 \leq$ $j<k$, we make an assumption about $g_{k} \in M \mathcal{K}_{m}\left(\left(A-\rho_{k} B\right) M, r_{k}\right)$ as follows.

Assumption 4.5. Let $g_{k}=M p_{m-1}^{(k)}\left(\left(A-\rho_{k} B\right) M\right) r_{k}$ at step 6 of Algorithm 2, where $p_{m-1}^{(k)}$ is a polynomial of degree no greater than $m-1$ with real coefficients. With a sufficiently small $\theta_{0}=\angle\left(x_{0}, v_{1}\right)_{B}$, assume that for all $0 \leq j<k$,

$$
M p_{m-1}^{(k)}\left(\left(A-\rho_{k} B\right) M\right)=M p_{m-1}^{(j)}\left(\left(A-\rho_{j} B\right) M\right)+\mathcal{O}\left(\sin \theta_{j}\right)
$$

Remark. Assumption 4.5 holds trivially for LOPCG, i.e., Algorithm 2 with $m=1$, because $p_{m-1}^{(k)}(\cdot)=I$ for all $k$, so that $g_{k}=M r_{k}$ up to a scaling factor.

Lemma 8. Under Assumptions 4.3 and 4.5, $p_{2}^{T}\left(A-\rho\left(x_{2}\right) B\right) p_{0}=\mathcal{O}\left(\sin \theta_{0}\right)$.
From Lemmas 7 and 8, the search directions $\left\{p_{0}, p_{1}, p_{2}\right\}$ are pairwise approximately conjugate. In addition, since $x_{2} \in \operatorname{span}\left\{x_{1}, g_{1}, p_{0}\right\}=\operatorname{span}\left\{x_{0}, p_{0}, p_{1}\right\}=U_{2}$ is obtained from the RR projection, it is the global minimizer in $U_{2}$. Let $z^{*}$ be the global minimizer of $\rho(\cdot)$ in $U_{3}$. It follows from Theorem 6 that

$$
\begin{equation*}
\rho\left(x_{3}\right)-\rho\left(z^{*}\right) \leq \rho\left(x_{2}\right)-\rho\left(y^{*}\right)+\mathcal{O}\left(\sin \theta_{3}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right)=\mathcal{O}\left(\sin \theta_{3}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right) \tag{4.15}
\end{equation*}
$$

where $x_{2}=y^{*}$ is the global minimizer in $U_{2}$. We thus have the base case: $x_{3}$ is a global quasi minimizer in $U_{3}$.

The rest of our work is focused on the inductive step: assuming that the global quasi optimality is achieved at $x_{k}$, we want to show that the new $p_{k}$ is approximately conjugate to $\left\{p_{0}, \ldots, p_{k-1}\right\}$ such that the quasi optimality can be extended to $x_{k+1}$.

Lemma 9. For a given $k \geq 3$, assume that the iterate $x_{k}$ of Algorithm 2 achieves the global quasi optimality (Definition 4.2). Then $r_{k}=\left(A-\rho\left(x_{k}\right) B\right) x_{k}$ satisfies

$$
\begin{equation*}
\lim _{\theta_{0} \rightarrow 0} \cos \angle\left(r_{k}, W_{k}\right)=0 \tag{4.16}
\end{equation*}
$$

i.e., $r_{k}$ is asymptotically orthogonal to $W_{k}=\operatorname{span}\left\{p_{0}, \ldots, p_{k-1}\right\}$.

Note that $\left.\cos \angle\left(r_{k}, W_{k}\right)\right|_{\theta_{0}=0}$ is not defined, since Algorithm 2 with $x_{0}=v_{1}$ would not proceed. However, thanks to (4.16), it is reasonable for us to make the following assumption about the behavior $\cos \angle\left(r_{k}, W_{k}\right)$ near $\theta_{0}=0$.

Assumption 4.6. For a given step $k \geq 3$, suppose that the iterate $x_{k}$ of Algorithm 2 achieves global quasi optimality (4.7). We define $\left.\cos \angle\left(r_{k}, W_{k}\right)\right|_{\theta_{0}=0}=0$ and assume that $\cos \angle\left(r_{k}, W_{k}\right)$ is differentiable at $\theta_{0}=0$; that is, we assume that $\cos \angle\left(r_{k}, W_{k}\right) \leq$ $\mathcal{O}\left(\sin \theta_{0}\right)=\mathcal{O}\left(\sin \theta_{\ell}\right)(1 \leq \ell \leq k)$ for a sufficiently small $\theta_{0}$.

We note that Assumption 4.6 is not just presented for technical convenience but is consistent with our numerical experience, at least for relatively small $k$. Under such an assumption, we can establish the inductive step as follows.

Theorem 10. Assume that the $B$-normalized $\left\{p_{0}, \ldots, p_{k-1}\right\}$ of Algorithm 2 are approximately conjugate, and hence $x_{k}$ achieves global quasi optimality (4.7). Under Assumptions 4.3, 4.4, and 4.5, $\left\{p_{0}, \ldots, p_{k}\right\}$ are approximately conjugate.

Proof. To complete the proof, it is sufficient to show that $p_{k}$ is approximately conjugate to $p_{\ell}$ for all $0 \leq \ell \leq k-2$. For any $\ell, 0 \leq \ell \leq k-2$, note that $x_{\ell+1}=$ $\beta_{\ell(\ell+1)} x_{\ell}+\gamma_{\ell(\ell+1)} p_{\ell}$, where $\gamma_{\ell(\ell+1)}=\mathcal{O}\left(\sin \theta_{\ell}\right)$ and $\beta_{\ell(\ell+1)}=1-\mathcal{O}\left(\sin \theta_{\ell}\right)$. It follows that $p_{\ell}=\frac{1}{\gamma_{\ell(\ell+1)}}\left(x_{\ell+1}-\beta_{\ell(\ell+1)} x_{\ell}\right)$, and

$$
\begin{aligned}
& \left|p_{\ell}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{k}\right|=\frac{1}{\left|\gamma_{\ell(\ell+1)}\right|}\left|\left(x_{\ell+1}^{T}-\beta_{\ell(\ell+1)} x_{\ell}^{T}\right)\left(A-\rho\left(x_{k}\right) B\right) g_{k}\right| \\
& \left.=\frac{1}{\left|\gamma_{\ell(\ell+1)}\right|} \right\rvert\,\left[x_{\ell+1}^{T}\left(A-\rho\left(x_{\ell+1}\right) B\right)-\beta_{\ell(\ell+1)} x_{\ell}^{T}\left(A-\rho\left(x_{\ell}\right) B\right)\right] g_{k} \\
& +\left(\rho\left(x_{\ell+1}\right)-\rho\left(x_{k}\right)\right) x_{\ell+1}^{T} B g_{k}-\beta_{\ell(\ell+1)}\left(\rho\left(x_{\ell}\right)-\rho\left(x_{k}\right)\right) x_{\ell}^{T} B g_{k} \mid, \\
& v b=\frac{\left|\left(r_{\ell+1}^{T}-\beta_{\ell(\ell+1)} r_{\ell}^{T}\right) g_{k}+\mathcal{O}\left(\sin ^{2} \theta_{\ell+1}\right)\left\|B x_{\ell+1}\right\|\left\|g_{k}\right\|+\mathcal{O}\left(\sin ^{2} \theta_{\ell}\right)\left\|B x_{\ell}\right\|\left\|g_{k}\right\|\right|}{\left|\gamma_{\ell(\ell+1)}\right|} \\
& \leq \frac{\mathcal{O}\left(\sin ^{2} \theta_{\ell}\right)\left\|g_{k}\right\|}{\mathcal{O}\left(\sin \theta_{\ell}\right)}+\frac{\left|\left(r_{\ell+1}^{T}-\beta_{\ell(\ell+1)} r_{\ell}^{T}\right) M p_{m-1}^{(k)}\left(\left(A-\rho_{k} B\right) M\right) r_{k}\right|}{\left|\gamma_{\ell(\ell+1)}\right|} \\
& \left.=\mathcal{O}\left(\sin \theta_{\ell}\right)\left\|g_{k}\right\|+\frac{1}{\left|\gamma_{\ell(\ell+1)}\right|} \right\rvert\, r_{\ell+1}^{T}\left[p_{m-1}^{(\ell+1)}\left(M\left(A-\rho_{\ell+1} B\right)\right) M+\mathcal{O}\left(\sin \theta_{\ell+1}\right)\right] r_{k} \\
& -\beta_{\ell(\ell+1)} r_{\ell}^{T}\left[p_{m-1}^{(\ell)}\left(M\left(A-\rho_{\ell} B\right)\right) M+\mathcal{O}\left(\sin \theta_{\ell}\right)\right] r_{k} \mid .
\end{aligned}
$$

Since $p_{m-1}^{(\ell)}\left(M\left(A-\rho_{\ell} B\right)\right) M$ is symmetric, $r_{\ell} p_{m-1}^{(\ell)}\left(M\left(A-\rho_{\ell} B\right)\right) M=g_{\ell}^{T}$, and this holds similarly if $\ell$ is replaced with $\ell+1$. Also, since $x_{k}$ achieves global quasi optimality, by Assumption 4.6, $\left|\cos \angle\left(r_{k}, W_{k}\right)\right| \leq \mathcal{O}\left(\sin \theta_{\ell}\right)(0 \leq \ell \leq k)$. Since $g_{\ell}, g_{\ell+1} \in W_{k}$, we have $\left|\cos \angle\left(r_{k}, g_{\ell}\right)\right| \leq \mathcal{O}\left(\sin \theta_{\ell}\right)$ and $\left|\cos \angle\left(r_{k}, g_{\ell+1}\right)\right| \leq \mathcal{O}\left(\sin \theta_{\ell+1}\right)$. Note that $\left\|r_{\ell}\right\|=\mathcal{O}\left(\sin \theta_{\ell}\right)$, and $\left\|r_{k}\right\|=\mathcal{O}\left(\sin \theta_{k}\right)=\mathcal{O}\left(\sin \theta_{\ell+1}\right)=\mathcal{O}\left(\sin \theta_{\ell}\right)$. Hence,

$$
\begin{align*}
& \quad\left|p_{\ell}\left(A-\rho\left(x_{k}\right) B\right) g_{k}\right|  \tag{4.17}\\
& \leq\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)+\frac{\left|g_{\ell+1}^{T} r_{k}+r_{\ell+1}^{T} r_{k} \mathcal{O}\left(\sin \theta_{\ell+1}\right)-\beta_{\ell(\ell+1)} g_{\ell}^{T} r_{k}+r_{\ell}^{T} r_{k} \mathcal{O}\left(\sin \theta_{\ell}\right)\right|}{\mathcal{O}\left(\sin \theta_{\ell}\right)} \\
& \leq\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)+\frac{\left\|g_{\ell+1}\right\|\left\|r_{k}\right\|\left|\cos \angle\left(r_{k}, g_{\ell+1}\right)\right|+\left|\beta_{\ell(\ell+1)}\right|\left\|g_{\ell}\right\|\left\|r_{k}\right\|\left|\cos \angle\left(r_{k}, g_{\ell}\right)\right|}{\mathcal{O}\left(\sin \theta_{\ell}\right)} \\
& \quad+\left\|r_{\ell+1}\right\|\left\|r_{k}\right\| \mathcal{O}(1)+\left\|r_{\ell}\right\|\left\|r_{k}\right\| \mathcal{O}(1) \\
& \leq\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)+\frac{\left\|r_{k}\right\|\left(\left\|g_{\ell+1}\right\| \mathcal{O}\left(\sin \theta_{\ell+1}\right)+\left|\beta_{\ell(\ell+1)}\right|\left\|g_{\ell}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)\right)}{\mathcal{O}\left(\sin \theta_{\ell}\right)}+\mathcal{O}\left(\sin ^{2} \theta_{\ell}\right) \\
& \leq\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)+\left\|g_{\ell+1}\right\| \mathcal{O}\left(\sin \theta_{k}\right)+\left\|g_{\ell}\right\| \mathcal{O}\left(\sin \theta_{k}\right)+\mathcal{O}\left(\sin ^{2} \theta_{\ell}\right)=\mathcal{O}\left(\sin \theta_{\ell}\right) .
\end{align*}
$$

Moreover, it is not difficult to see that $\left|p_{\ell}\left(A-\rho\left(x_{k}\right) B\right) g_{k}\right|=\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)$. Recall that $\widetilde{p}_{k}=g_{k}-\frac{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}} p_{k-1}$ such that $\left\|\widetilde{p}_{k}\right\|$ is proportional to $\left\|g_{k}\right\|$. The normalized search direction is $p_{k}=\frac{\eta_{k}}{\left\|g_{k}\right\|} \widetilde{p}_{k}$, where $\eta_{k}$ is chosen such that $\left\|p_{k}\right\|_{B}=1$. Also, since $0 \leq \ell \leq k-2, p_{\ell}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}=\mathcal{O}\left(\sin \theta_{\ell}\right)$, and

$$
\begin{align*}
& p_{\ell}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{k}=\frac{\eta_{k}}{\left\|g_{k}\right\|} p_{\ell}^{T}\left(A-\rho\left(x_{k}\right) B\right) \widetilde{p}_{k}  \tag{4.18}\\
&= \frac{\eta_{k}}{\left\|g_{k}\right\|} p_{\ell}^{T}\left(A-\rho\left(x_{k}\right) B\right)\left(g_{k}-\frac{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}} p_{k-1}\right) \\
&= \frac{\eta_{k}}{\left\|g_{k}\right\|}\left\{p_{\ell}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{k}-\frac{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}}\right. \\
&\left.\quad \times\left[p_{\ell}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}+\left(\rho\left(x_{k-1}\right)-\rho\left(x_{k}\right)\right) p_{\ell}^{T} B p_{k-1}\right]\right\} \\
&= \frac{\eta_{k}}{\left\|g_{k}\right\|}\left\{\left\|g_{k}\right\| \mathcal{O}\left(\sin \theta_{\ell}\right)+\mathcal{O}\left(\left\|g_{k}\right\|\right)\left[\mathcal{O}\left(\sin \theta_{\ell}\right)+\mathcal{O}\left(\sin ^{2} \theta_{k-1}\right)\right]\right\}=\mathcal{O}\left(\sin \theta_{\ell}\right) .
\end{align*}
$$

Finally, recall from Lemma 7 that $p_{k-1}$ and $p_{k}$ are approximately conjugate. It follows that $\left\{p_{0}, \ldots, p_{k}\right\}$ are pairwise approximately conjugate.

With all the above work, we can finally present the major theorem of this section on the global quasi optimality of all iterates generated by Algorithm 2 .

Theorem 11. Let $\left\{x_{k}\right\}$ be the iterates of Algorithm $2, r_{k}=\left(A-\rho\left(x_{k}\right) B\right) x_{k}$ the residual, and $\left\{p_{k}\right\}$ the $B$-normalized search directions. Suppose that Assumptions 4.3, 4.4, and 4.5 hold. In Theorem 6 , let $y^{*}$ and $z^{*}$ be the global minimizer in $U_{k}$ and $U_{k+1}$, respectively, and $C_{k+1}$ be the constant for which $\rho\left(x_{k+1}\right)-\rho\left(z^{*}\right) \leq \rho\left(x_{k}\right)-\rho\left(y^{*}\right)+$ $C_{k+1} \sin \theta_{k+1} \sin ^{2} \theta_{0}$ if $\left\{p_{0}, \ldots, p_{k}\right\}$ are approximately conjugate. Then for any given $k \geq 3, \rho\left(x_{k}\right)-\rho\left(y^{*}\right) \leq \sum_{j=3}^{k} C_{j} \sin \theta_{j} \sin ^{2} \theta_{0}$, and therefore,

$$
\begin{equation*}
\text { 9) } \lim _{\theta_{0} \rightarrow 0} \frac{\rho\left(x_{k}\right)-\rho\left(y^{*}\right)}{\rho\left(x_{k}\right)-\lambda_{1}} \leq \lim _{\theta_{0} \rightarrow 0} \frac{\sum_{j=3}^{k} C_{j} \sin \theta_{j} \sin ^{2} \theta_{0}}{\underline{\xi}^{k}\left(\rho\left(f_{0}\right)-\lambda_{1}\right) \sin ^{2} \theta_{0}} \leq \lim _{\theta_{0} \rightarrow 0} \frac{\left(\underline{\xi}^{-1}\right)^{k} \sum_{j=3}^{k} C_{j} \sin \theta_{j}}{\lambda_{2}-\lambda_{1}}=0 \text {. } \tag{4.19}
\end{equation*}
$$

Proof. The proof is based on mathematical induction. The base case has been established in Lemmas 7 and 8 together with (4.15). The inductive step is completed in Theorem 10. Then, by Theorem 6, the conclusion holds.

Remark. Interestingly, (4.19) suggests that for a larger $k$, the relative difference between $\rho\left(x_{k}\right)$ (locally optimized $\rho(\cdot)$ over $\operatorname{span}\left\{x_{k-1}, g_{k-1}, p_{k-2}\right\}$ ) and $\rho\left(y^{*}\right)$ (globally optimized $\rho(\cdot)$ over $\left.U_{k}=\operatorname{span}\left\{x_{0}\right\}+\operatorname{span}\left\{g_{0}, \ldots, g_{k-1}\right\}\right)$ tends to grow larger. The upper bound on this relative difference scales like a linear function of $k$ (assuming that $\left|C_{j} \sin \theta_{j}\right| \leq C$ for all $j \geq 3$ ) multiplied by an exponentially increasing function of $k$. Nevertheless, for a given outer iteration $k \geq 3$, Theorem 11 shows that Algorithm 2 iterate $x_{k} \in U_{k}$ is almost as good as the corresponding global minimizer $y^{*}$ if $\theta_{0}$ is sufficiently small. As a result, Algorithm 2 would actually converge considerably faster than Theorem 1 suggests. This theorem also provides insight into the performance of LOPCG, which is an instance of Algorithm 2, and the block extension LOBPCG.

Remark. A different viewpoint is followed in [36] but with qualitatively similar results. Consider $k$ steps of CG solving the JD correction equation starting from the Ritz vector at the $i$ th iteration of Lanczos. Then, the distance between the Ritz vector of Lanczos at iteration $i+k$ and the CG solution after $k$ steps is bounded by $O\left(\left|\rho_{i}-\rho_{i+k}\right|\right)$. The locally optimal restarting approximates the $A$-norm minimization of CG over the entire space, so if Lanczos converges slowly or if $k$ is small, +K restarting will yield Ritz vectors close to the unrestarted Lanczos. If $\rho_{i}$ and $\rho_{i+k}$ are far, then convergence is already fast and the use of +K is not needed.
5. Experiments. First we show that TRPL+K does indeed achieve quasi optimality and then we investigate the effect of the maximum basis size $q$ and the number of previous retained vectors $l$ on the performance. Then we present an extensive set of experiments comparing TRPL +K against other methods on a variety of problems.

The matrix set is chosen to overlap with other papers in the literature. All matrices can be reproduced from these papers or downloaded from the SuiteSparse Matrix collection (formerly at the University of Florida). Table 5.1 lists some basic properties of these matrices. Matrices finan512 and Plate33K_A0 have larger gap ratios so they are relatively easy problems. Matrices Trefethen_20k and cfd1 are moderately hard problems, while matrices 1138bus and or56f are the hardest problems. The effectiveness of a method is typically manifested on harder problems.

We compare TRPL+K against the following eigenmethods: unrestarted GD (as a representative of unrestarted Lanczos), TRLan [43], GD+K [34, 36], LOBPCG [17], and EigIFP [11]. To allow for easier comparisons, we use only MATLAB implementations. We adopt the publicly available implementations for LOBPCG and EigIFP

Table 5.1
Properties of the test matrices for standard eigenvalue problems. Trefethen_20k is Tre $20 k$, and Plate33K_A0 is PlaA0.

| Matrix | Order | $\mathrm{nnz}(\mathrm{A})$ | $\kappa(A)$ | Source |
| :---: | :---: | :---: | :---: | :---: |
| 1138bus | 1138 | 4054 | 1.2 E 7 | MM |
| or56f | 9000 | 2890150 | 6.6 E 7 | Yang |
| Tre20k | 20000 | 554466 | 2.9 E 5 | UF |
| PlaA0 | 39366 | 914116 | 6.8 E 9 | FEAP |
| cfd1 | 70656 | 1825580 | 1.3 E 6 | UF |
| finan512 | 74752 | 596992 | 9.8 E 1 | UF |

while we provide implementations for TRPL+K, TRLan, and GD+K. All methods employ the same stopping criterion satisfying

$$
\begin{equation*}
\left\|A u_{i}-\theta_{i} u_{i}\right\| \leq\|A\|_{F} \delta_{u s e r} \tag{5.1}
\end{equation*}
$$

where $\|A\|_{F}$ is the Frobenius norm of $A$ and $\delta_{\text {user }}$ is a user-specified stopping tolerance. All experiments use $\delta=1 \mathrm{E}-14$. All methods start with the same initial guess, $\operatorname{rand}(n, k)$, with fixed seed number (12). We set the maximum number of restarts to 5000 for all methods. Since we focus on methods with limited memory, we set $\operatorname{maxBasisSize}=18$, minRestartSize $=8$ for TRPL+K, TRLan, and GD+K. Since EigIFP restarts with a single vector, we set maxBasisSize $=18$ for its inner iteration. For LOBPCG, the method always uses $\operatorname{maxBasisSize}=3 p$. For other parameters we follow the defaults suggested in each code. We seek $p=1,5$, and 10 algebraically smallest eigenpairs for both stardard and generalized eigenvalue problems. We employ soft locking since the desired number of eigenpairs is small. All computations are carried out on an Apple MacBook Pro with Intel Core i7 processors at 2.2 GHz for a total of 4 cores and 16 GB of memory running the Mac Unix operating system. We use MATLAB 2016a with machine precision $\epsilon=2.2 \mathrm{E}-16$. We compare both the number of matrix-vector products (reported as "MV" in the tables) and runtime in seconds (reported as "Sec").
5.1. Quasi optimality with + K. Quasi optimality as defined in (4.7) implies that as the initial guess becomes better, the relative difference between the vector iterates of TRPL +K and the unrestarted method tends to zero. Figure 5.1 demonstrates that TRPL +K achieves this quasi optimality on two sample matrices. A standard


FIG. 5.1. Confirming the quasi optimality of TRPL+1 with no inner iterations by showing that the ratio (4.7) converges to zero as the angle of the initial guess to the required eigenvector decreases. We plot this for the first 100 iterations.
eigenproblem is solved without preconditioning. TRPL+K uses only one inner iteration (i.e., equivalent to LOPCG) because it is easier to compare step by step to the unrestarted method. The remark after Theorem 11 suggests that the ratio in (4.7) increases with the number of iterations, which we observe in the plots. Therefore, we plot only the first 100 iterations. What is important, however, is that for any given iteration, the ratio decreases as we make the initial guess better. Therefore, as TRPL+K converges asymptotically, it increasingly matches the unrestarted method.
5.2. The effects of maxBasisSize $\boldsymbol{q}$ and maxPrevSize $\boldsymbol{l}$. First we investigate the effects of varying $q$ from 8 to 512 on the performance of TRPL+K (without a preconditioner) and TRLan compared against unrestarted Lanczos, with $l=1$ and $p=1$. We choose three hard problems where the differences are pronounced. Figure 5.2 shows that as the maximum basis size increases, both TRPL+K and TRLan become similar to the unrestarted Lanczos. However, while TRLan requires the increased basis to significantly improve convergence, TRPL+K achieves very good or even close to optimal convergence with very small basis size. The slight increase of matvecs in TRPL +K with very large basis sizes may be attributed to the more targeted expansion of the subspace using the residual instead of the Lanczos vectors.

We then vary the number of previous vectors $l$ from 0 to 5 to study its impact on the convergence of TRPL +K (without a preconditioner), seeking $p=5$ smallest eigenvalues. When $l=0$ and without preconditioning, TRPL +K reduces to TRLan. Figure 5.3 shows the similar trend for all three cases: a significant reduction in matvecs with $l=1$, while $l>1$ results in a slight deterioration in convergence. This is qualitatively similar to earlier observations for GD +K , where $l>1$ is beneficial only with a block method. This implies that we obtain all the benefits of the +K technique with minimal overhead.


Fig. 5.2. The effect of varying the maximum basis size $q$ on the convergence of $T R P L+K$ and TRLan. Seeking $p=1$ smallest eigenvalue without preconditioning.


Fig. 5.3. The effect of varying the number of previous vectors $l$ on the convergence of TRPL+K when seeking $p=5$ smallest eigenvalues without preconditioning.
5.3. Without preconditioning. We compare against unrestarted Lanczos and other methods without a preconditioner to show the near-optimal performance of TRPL +K under limited memory. Figure 5.4 shows the results on the easy case finan512 and on the hard case cfd1, seeking $p=1$ eigenvalue. For both cases TRPL+K and GD +K can achieve almost identical performance as unrestarted Lanczos, which is consistent with Theorem 11. GD+K requires slightly fewer iterations because its new directions come from the most recent $R R$ over the entire space, although the cost per iteration is higher. Compared to LOBPCG and EigIFP, TRPL+K is significantly faster because it uses thick restarting and a Krylov subspace (compared to LOBPCG) and because it uses thick and locally optimal restarting (compared to EigIFP). The difference between TRPL +K and TRLan is relatively small for the easy case but becomes significant with hard problems. Figure 5.5 shows similar results when computing 5 eigenvalues. To make it easier to see, we only show graphs for TRPL+K, $\mathrm{GD}+\mathrm{K}$, and TRLan. As before, TRPL+K and GD+K are competitive and both of them substantially outperform the performance of TRLan.

Tables 5.2 and 5.3 summarize results from all methods seeking 1 and 5 eigenvalues, respectively. TRPL +K converges faster than all other methods in terms of matrixvector products except for GD+K. In fact, the harder the problem the more significant


FIG. 5.4. Compare all methods to the optimal convergence of unrestarted Lanczos when seeking 1 smallest eigenvalue without preconditioning.


Fig. 5.5. Compare $T R P L+K$ and $G D+K$ to the optimal convergence of unrestarted Lanczos when seeking 5 smallest eigenvalue without preconditioning.

TABLE 5.2
Comparing all methods for 1 smallest eigenvalue without preconditioning.

| Method: | TRPL+K |  | TRLan |  |  |  |  |  |  |  |  | GD+K |  | LOBPCG |  | EIGIFP |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Matrix | MV | Sec | MV | Sec | MV | Sec | MV | Sec | MV | Sec |  |  |  |  |  |  |  |
| 1138BUS | 4778 | $\mathbf{0 . 9}$ | 48718 | 8.4 | $\mathbf{3 6 6 2}$ | 1.8 | 15498 | 6.4 | 11107 | 1.2 |  |  |  |  |  |  |  |
| or56f | 4908 | $\mathbf{2 0 . 2}$ | 64058 | 257.5 | $\mathbf{4 0 2 4}$ | 21.5 | 20365 | 82.1 | 19117 | 72.3 |  |  |  |  |  |  |  |
| Tre20k | 2208 | $\mathbf{7 . 1}$ | 12718 | 36.4 | $\mathbf{2 0 5 8}$ | 10.3 | 3482 | 8.6 | 3799 | 8.0 |  |  |  |  |  |  |  |
| PlaA0 | 638 | $\mathbf{3 . 5}$ | 1168 | 5.5 | $\mathbf{5 5 6}$ | 6.4 | 1583 | 5.4 | 1207 | 3.9 |  |  |  |  |  |  |  |
| cfd1 | 2358 | $\mathbf{2 1 . 7}$ | 13008 | 106.7 | $\mathbf{1 9 7 0}$ | 44.9 | 5647 | 33.7 | 4807 | 29.5 |  |  |  |  |  |  |  |
| finan512 | 358 | 6.0 | 428 | 7.3 | $\mathbf{3 2 2}$ | 6.7 | 1028 | 4.8 | 685 | $\mathbf{3 . 3}$ |  |  |  |  |  |  |  |

Table 5.3
Comparing all methods for 5 smallest eigenvalue without preconditioning.

| Method: | TRPL+K |  | TRLan |  | GD+K |  |  | LOBPCG |  | EIGIFP |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | :---: |
| Matrix | MV | Sec | MV | Sec | MV | Sec | MV | Sec | MV | Sec |  |
| 1138BUS | 26888 | $\mathbf{5 . 1}$ | 414518 | 74.9 | $\mathbf{2 1 3 1 8}$ | 11.2 | 374915 | 56.1 | 68729 | 7.9 |  |
| or56f | 13218 | $\mathbf{5 5 . 4}$ | 115098 | 507.7 | $\mathbf{1 1 1 6 9}$ | 61.2 | 40940 | 65.5 | 68801 | 260.6 |  |
| Tre20k | 6158 | $\mathbf{1 8 . 1}$ | 23358 | 65.8 | $\mathbf{5 5 2 0}$ | 35.4 | 14335 | 21.8 | 25277 | 55.0 |  |
| PlaA0 | 1858 | $\mathbf{1 0 . 0}$ | 2728 | 12.7 | $\mathbf{1 8 1 2}$ | 23.3 | 12545 | 22.1 | 4235 | 14.3 |  |
| cfd1 | 8558 | $\mathbf{8 5 . 8}$ | 43648 | 386.6 | $\mathbf{7 7 4 7}$ | 212.5 | 70170 | 458.6 | 20993 | 132.3 |  |
| finan512 | 1668 | 27.2 | 2838 | 44.3 | $\mathbf{1 5 0 1}$ | 43.5 | 8830 | 35.2 | 3551 | $\mathbf{1 8 . 2}$ |  |

the gains. It is also interesting to observe the differences between the methods which are due mainly to the use of different algorithmic components. Among them, the locally optimal restarting has the biggest effect, followed by a larger Krylov space, while a combination of all three components in TRPL +K (and GD+K) are clearly the most beneficial while being computationally economical. Also, TRPL+K requires about $20 \%$ more matrix-vector products than GD+K but it is $35 \%$ faster in runtime, making it the method of choice when the matrix-vector operator is inexpensive. This is because TRPL +K only needs to perform the RR procedure once every outer iteration while GD+K needs to apply the RR procedure for every matrix-vector multiplication.
5.4. With preconditioning. The above results emphasize the need for preconditioning, especially for problems with highly clustered eigenvalues. In this case TRLan cannot be used. In our experiments, we use the MATLAB ILU factorization on $A$ with parameters 'type $=$ nofill'. We then compare TRPL +k against other methods with the constructed preconditioner for finding 1 and 5 smallest eigenvalues. The results are shown in Tables 5.4 and 5.5, respectively. TRPL +K is again the fastest or close to the fastest method. We note that with a good preconditioner the number of iterations decreases and thus the differences between methods are smaller. We also note that the added cost of the preconditioner per iteration is similar to having a more expensive operator which favors $\mathrm{GD}+\mathrm{K}$ also in runtime.

TABLE 5.4
Comparing all methods for 1 smallest eigenvalue with preconditioning.

| Method: | TRPL+K |  | GD+K |  | LOBPCG |  | EIGIFP |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| Matrix | MV | Sec | MV | Sec | MV | Sec | MV | Sec |
| 1138BUS | 328 | $\mathbf{0 . 2}$ | $\mathbf{1 7 4}$ | $\mathbf{0 . 2}$ | 474 | 0.3 | 397 | $\mathbf{0 . 2}$ |
| or56f | 158 | 1.5 | $\mathbf{8 1}$ | $\mathbf{0 . 8}$ | 335 | 2.7 | 325 | 2.7 |
| Tre20k | 38 | 0.2 | 17 | $\mathbf{0 . 1}$ | $\mathbf{1 0}$ | $\mathbf{0 . 1}$ | 55 | 0.3 |
| cfd1 | 838 | 20.0 | $\mathbf{5 6 2}$ | $\mathbf{1 4 . 0}$ | 2519 | 26.5 | 1279 | 17.2 |
| finan512 | 98 | 1.9 | $\mathbf{6 9}$ | 1.5 | 144 | $\mathbf{1 . 0}$ | 127 | 1.4 |

TABLE 5.5
Comparing all methods for 5 smallest eigenvalues with preconditioning.

| Method: | TRPL+K |  | GD+K |  | LOBPCG |  | EIGIFP |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| Matrix | MV | Sec | MV | Sec | MV | Sec | MV | Sec |
| 1138BUS | 1128 | $\mathbf{0 . 3}$ | $\mathbf{7 7 9}$ | 0.5 | 2120 | 0.5 | 2795 | 0.7 |
| or56f | 368 | 3.3 | $\mathbf{2 2 3}$ | 2.4 | 375 | $\mathbf{2 . 3}$ | 1175 | 10.1 |
| Tre20k | 118 | 0.7 | $\mathbf{4 5}$ | 0.4 | 115 | $\mathbf{0 . 3}$ | 203 | 1.0 |
| cfd1 | 2598 | 67.1 | $\mathbf{1 8 6 8}$ | $\mathbf{5 6 . 7}$ | 12900 | 89.7 | 5135 | 72.8 |
| finan512 | 378 | 8.6 | $\mathbf{2 8 9}$ | 9.0 | 1225 | $\mathbf{5 . 2}$ | 779 | 8.5 |

Table 5.6
Properties of the test matrices for generalized eigenvalue problems.

| Matrix | Order | $\mathrm{nnz}(\mathrm{A})$ | $\kappa(A)$ | Source |
| :---: | :---: | :---: | :---: | :---: |
| bcsstk23 | 3134 | 45178 | $6.9 \mathrm{e}+12$ | MM |
| bcsstm23 | 3134 | 3134 | $9.4 \mathrm{e}+08$ | MM |
| bcsstk24 | 3562 | 159910 | $6.3 \mathrm{e}+11$ | MM |
| bcsstm24 | 3562 | 3562 | $1.8 \mathrm{e}+13$ | MM |
| bcsstk25 | 15439 | 252241 | $1.2 \mathrm{e}+13$ | MM |
| bcsstm25 | 15439 | 15439 | $6.0 \mathrm{e}+09$ | MM |

TABLE 5.7
Comparing all methods for computing 1,5 , and 10 smallest eigenvalues with preconditioning (numbers in parenthesis). Each pencil bcsstkmXY has matrix pairs (bcsstkXY, bcsstmXY), where $X Y \in\{23,24,25\}$.

| Method: | TRPL+K |  | LOBPCG |  | EIGIFP |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| Matrix | MV | Sec | MV | Sec | MV | Sec |
| bcsstkm23(1) | $\mathbf{6 0}$ | $\mathbf{0 . 1}$ | 134 | 0.2 | 127 | 0.2 |
| bcsstkm24(1) | $\mathbf{2 9}$ | $\mathbf{0 . 1}$ | 55 | $\mathbf{0 . 1}$ | 73 | 0.2 |
| bcsstkm25(1) | $\mathbf{2 8 1 6}$ | $\mathbf{2 9 . 1}$ | 28355 | 178.9 | 9721 | 66.2 |
| bcsstkm23(5) | $\mathbf{1 5 0}$ | 0.4 | 190 | $\mathbf{0 . 3}$ | 419 | 0.7 |
| bcsstkm24(5) | $\mathbf{1 2 0}$ | $\mathbf{0 . 6}$ | 1020 | 3.6 | 257 | 0.8 |
| bcsstkm25(5) | $\mathbf{3 8 3 0}$ | $\mathbf{3 9 . 6}$ | 39895 | 89.1 | 14423 | 97.6 |
| bcsstkm23(10) | $\mathbf{2 1 0}$ | 0.4 | 240 | $\mathbf{0 . 2}$ | 820 | 1.0 |
| bcsstkm24(10) | $\mathbf{2 1 0}$ | 0.9 | 270 | $\mathbf{0 . 6}$ | 496 | 1.4 |
| bcsstkm25(10) | $\mathbf{5 1 2 0}$ | $\mathbf{5 2 . 9}$ | 45820 | 63.7 | 20494 | 141.7 |

5.5. Generalized eigenvalue problem. We perform some sample experiments on generalized eigenvalue problems, comparing the proposed method against LOBPCG and EigIFP. As shown in Table 5.6, the condition numbers of these problems are quite large, making preconditioning necessary to accelerate the convergence. In this experiment, we use incomplete LDL factorization [12] on $A$ with droptol $=1 \mathrm{e}-6$, $1 \mathrm{e}-8$ for bcsstkm23 and bcsstkm24, respectively, and the MATLAB LDL factorization on $A$ with parameter THRESH $=0.5$ for bcsstkm25. We then use the preconditioner to find 1,5 , and 10 smallest eigenvalues. As shown in Table 5.7 , TRPL+K significantly outperforms other methods in terms of the number of matrix-vector operations for these cases. Note that LOGPCG is quite efficient in terms of runtime, thanks to its efficiently implemented block operations in MATLAB, especially for small matrices that fit in cache. A block TRPL+K is also possible, which is left for future work.
6. Conclusions and future work. We presented a new near-optimal eigenmethod, thick-restart preconditioned Lanczos +K method (TRPL +K ), which is based on three key algorithmic components: thick restarting, locally optimal restarting, and the ability to build a preconditioned Krylov space. We provided a proof of an asymptotic global quasi optimality of the proposed method and provided insights on the
near-optimal performance of a group of methods that employ locally optimal restarting. Our extensive experiments demonstrate that TRPL+K either outperforms or matches other state-of-the-art methods in terms of both number of matrix-vector operations and computational time.

An interesting future direction is to extend this approach to the Lanczos bidiagonalization method for singular value problems.

## Appendix A.

Proof of Lemma 3. By definition, the optimal step size $\alpha^{*}$ is

$$
\begin{equation*}
\alpha^{*}=\arg \min _{\alpha>0} \rho(x+\alpha p)=\arg \min _{\alpha>0} \frac{(x+\alpha p)^{T} A(x+\alpha p)}{(x+\alpha p)^{T} B(x+\alpha p)} \tag{A.1}
\end{equation*}
$$

which can be found by letting $\frac{d}{d \alpha} \frac{1}{2} \rho(x+\alpha p)=0$. Using the quotient rule for differentiation, with some algebraic work, we have

$$
\begin{equation*}
\frac{d}{d \alpha} \frac{\rho(x+\alpha p)}{2}=\frac{a(x, p) \alpha^{2}+b(x, p) \alpha+c(x, p)}{\left[(x+\alpha p)^{T} B(x+\alpha p)\right]^{2}} \tag{A.2}
\end{equation*}
$$

where $a(x, p), b(x, p)$, and $c(x, p)$ are given in (4.12). For the sake of simplicity, we refer to these coefficients as $a, b$, and $c$, respectively, when there is no danger of confusion. Depending on the sign of $a$, there are three cases for the solutions of $\frac{d}{d \alpha} \rho(x+\alpha p)=0$ :

1. $a=0$. Obviously, there is a unique solution $\alpha^{*}=-\frac{c}{b}>0$.
2. $a>0$. Since $c<0$, there is a unique positive root $\alpha^{*}=\frac{1}{2 a}\left(-b+\sqrt{b^{2}-4 a c}\right)$. By assumption, $b \geq\|x\|_{B}^{2}\|p\|_{B}^{2} \delta>0$, and $|c| \leq\|x\|_{B}^{2}\|p\|\|A x-\rho(x) B x\|=$ $\mathcal{O}(\sin \theta)$ for sufficiently small $\theta$. By the Taylor expansion of $\sqrt{1-t}$ for $|t|<1$,

$$
\alpha^{*}=\frac{1}{2 a}\left(-b+b \sqrt{1-\frac{4 a c}{b^{2}}}\right)=-\frac{c}{b}\left(1+\frac{a c}{b^{2}}+\frac{2 a^{2} c^{2}}{b^{4}}+\cdots\right)
$$

which is slightly smaller than $\alpha^{*}$ in case 1 because $a>0$ and $c<0$.
3. $a<0$. In this case, there must be two distinct solutions $0<\alpha_{1}^{*}<\alpha_{2}^{*}$ such that $\frac{d}{d \alpha} \rho(x+\alpha p)<0$ for $0 \leq \alpha<\alpha_{1}^{*}$ or $\alpha>\alpha_{2}^{*}$, and $\frac{d}{d \alpha} \rho(x+\alpha p)>0$ for $\alpha_{1}^{*}<\alpha<\alpha_{2}^{*}$. In fact, if there is no solution or only one repeated solution, then $\frac{d}{d \alpha} \rho(x+\alpha p) \leq 0$ for all $\alpha \geq 0$, and hence $\rho(p)=\lim _{\alpha \rightarrow+\infty} \rho(x+\alpha p) \leq \rho(x)$, contradicting our assumption that $\rho(p)-\rho(x) \geq \delta>0$. Given these intervals of monotonicity, $\rho\left(x+\alpha_{2}^{*} p\right)>\lim _{\alpha \rightarrow+\infty} \rho(x+\alpha p)=\rho(p) \geq \rho(x)+\delta$. Hence the minimizer of $\rho(x+\alpha p)$ is achieved at $\alpha_{1}^{*}$ with the same expression as in case 2 . The optimal step size is slightly greater than $\alpha^{*}=-c / b$ in case 1 as $a, c<0$. We still refer to $\alpha_{1}^{*}$ as $\alpha^{*}$ for notation consistency.
In summary, $\rho(x+\alpha p)$ decreases monotonically on $\left(0, \alpha^{*}\right)$, then increases on $\left(\alpha^{*},+\infty\right)($ cases 1 and 2$)$ or on $\left(\alpha^{*}, \alpha_{2}^{*}\right)($ case 3$)$. The optimal $\alpha^{*}$ has a closed form. $\square$

Proof of Lemma 4. We note that the denominator of $\frac{d}{d \alpha} \rho(x+\alpha p)$, namely, $(x+$ $\alpha p)^{T} B(x+\alpha p)=\|p\|_{B}^{2} \alpha^{2}+2 p^{T} B x \alpha+\|x\|_{B}^{2}>0$ for all $\alpha$, and it is a quadratic with positive quadratic term coefficient. Hence, $\max _{0 \leq \alpha \leq \alpha^{*}}(x+\alpha p)^{T} B(x+\alpha p)=$ $\max _{\alpha=\left\{0, \alpha^{*}\right\}}(x+\alpha p)^{T} B(x+\alpha p)=\max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}$, where $\alpha^{*}=-\frac{c}{b}$ or $-\frac{c}{b}\left(1+\frac{a c}{b^{2}}+\ldots\right)$; see Lemma 3. For sufficiently small $\theta=\angle\left(x, v_{1}\right)_{B}$, since $|c| \leq$ $\mathcal{O}(\sin \theta)$, there is a small constant $\eta \in(0,1)$ independent of $\theta$, such that

$$
\begin{aligned}
& \rho\left(x+\alpha^{*} p\right)-\rho(x)=\int_{0}^{\alpha^{*}} \frac{d}{d \alpha} \rho(x+\alpha p) d \alpha=\int_{0}^{\alpha^{*}} \frac{a \alpha^{2}+b \alpha+c}{(x+\alpha p)^{T} B(x+\alpha p)} d \alpha \\
\leq & \frac{1}{\max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \int_{0}^{\alpha^{*}}\left(a \alpha^{2}+b \alpha+c\right) d \alpha=\frac{\frac{1}{3} a\left(\alpha^{*}\right)^{3}+\frac{1}{2} b\left(\alpha^{*}\right)^{2}+c \alpha^{*}}{\max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \\
= & \frac{-\frac{c^{2}}{2 b}-\frac{a c^{3}}{3 b^{3}}+\mathcal{O}\left(c^{4}\right)}{\max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \leq-\frac{c^{2}(1-\eta)}{2 b \max \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}},
\end{aligned}
$$

where $b \geq\|x\|_{B}^{2}\|p\|_{B}^{2} \delta>0$ is bounded away from zero; see (4.12).
To establish the lower bound, note that our assumption that $\rho(p)-\rho(x) \geq \delta>0$ for a constant $\delta$ independent of $\theta=\angle\left(x, v_{1}\right)_{B}$ means that there exists a $\tilde{\delta}>0$ independent of $\theta$, such that $\angle(p, x) \geq \tilde{\delta}$. Therefore, $\left\|x+\alpha^{*} p\right\|_{B}$ is bounded away from zero. Similarly, for sufficiently small $\theta$, we also have

$$
\begin{align*}
& \rho\left(x+\alpha^{*} p\right)-\rho(x) \geq \frac{1}{\min \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \int_{0}^{\alpha^{*}}\left(a \alpha^{2}+b \alpha+c\right) d \alpha  \tag{A.3}\\
= & \frac{-\frac{c^{2}}{2 b}-\frac{a c^{3}}{3 b^{3}}+\mathcal{O}\left(c^{4}\right)}{\min \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}} \geq-\frac{c^{2}(1+\eta)}{\min \left\{\|x\|_{B}^{2},\left\|x+\alpha^{*} p\right\|_{B}^{2}\right\}},
\end{align*}
$$

where $c=\|x\|_{B}^{2} p^{T}(A x-\rho(x) B x)=\mathcal{O}(\sin \theta) \mathcal{O}(\cos \angle(p, \nabla \rho(x)))$ from (4.12).
Proof of Lemma 5. First, note that one can always choose $\beta_{j k}, \gamma_{j k}>0$ in (4.13). If a decomposition has negative $\beta_{j k}$ or $\gamma_{j k}$, simply replace $x_{j}$ with $-x_{j}$ or $g_{j k}$ with $-g_{j k}$.

Given the decomposition (4.13), the normalization condition $\left\|x_{k}\right\|_{B}=1$ leads to

$$
\begin{aligned}
& \left\|x_{k}\right\|_{B}^{2}=\left(\beta_{j k} x_{j}+\gamma_{j k} g_{j k}\right)^{T} B\left(\beta_{j k} x_{j}+\gamma_{j k} g_{j k}\right) \\
(\mathrm{A} .4)= & \left\|x_{j}\right\|_{B}^{2} \beta_{j k}^{2}+2\left(g_{j k}^{T} B x_{j}\right) \gamma_{j k} \beta_{j k}+\left\|g_{j k}\right\|_{B}^{2} \gamma_{j k}^{2}=\beta_{j k}^{2}+2 \mu_{j k} \gamma_{j k} \beta_{j k}+\gamma_{j k}^{2} \\
= & \left(\gamma_{j k}+\mu_{j k} \beta_{j k}\right)^{2}+\left(1-\mu_{j k}^{2}\right) \beta_{j k}^{2}=1,
\end{aligned}
$$

where $\mu_{j k}=g_{j k}^{T} B x_{j}=\left(x_{j}, g_{j k}\right)_{B} \in[-1,1]$ because $\left\|x_{j}\right\|_{B}=\left\|g_{j k}\right\|_{B}=1$.
Let $r_{j}=A x_{j}-\rho\left(x_{j}\right) B x_{j}$ be the eigenresidual associated with $x_{j}$, and define $d_{j k} \equiv \rho\left(g_{j k}\right)-\rho\left(x_{j}\right)$. The Rayleigh quotient of $x_{k}$ is therefore
(A.5) $\rho\left(x_{k}\right)=\frac{x_{k}^{T} A x_{k}}{x_{k}^{T} B x_{k}}=x_{k}^{T} A x_{k}=\left(\beta_{j k} x_{j}+\gamma_{j k} g_{j k}\right)^{T} A\left(\beta_{j k} x_{j}+\gamma_{j k} g_{j k}\right)$

$$
\begin{aligned}
& =\left(x_{j}^{T} A x_{j}\right) \beta_{j k}^{2}+2\left(g_{j k}^{T} A x_{j}\right) \gamma_{j k} \beta_{j k}+\left(g_{j k}^{T} A g_{j k}\right) \gamma_{j k}^{2} \\
& =\rho\left(x_{j}\right) \beta_{j k}^{2}+2\left[g_{j k}^{T}\left(r_{j}+\rho\left(x_{j}\right) B x_{j}\right)\right] \gamma_{j k} \beta_{j k}+\rho\left(g_{j k}\right) \gamma_{j k}^{2} \quad\left(\text { as } A x_{j}=r_{j}+\rho\left(x_{j}\right) B x_{j}\right) \\
& =\rho\left(x_{j}\right) \beta_{j k}^{2}+2 \rho\left(x_{j}\right)\left(g_{j k}^{T} B x_{j}\right) \gamma_{j k} \beta_{j k}+2\left(g_{j k}^{T} r_{j}\right) \gamma_{j k} \beta_{j k}+\left(\rho\left(x_{j}\right)+d_{j k}\right) \gamma_{j k}^{2} \\
& =\rho\left(x_{j}\right)\left(\beta_{j k}^{2}+2 \mu_{j k} \gamma_{j k} \beta_{j k}+\gamma_{j k}^{2}\right)+2\left(g_{j k}^{T} r_{j}\right) \gamma_{j k} \beta_{j k}+d_{j k} \gamma_{j k}^{2} \\
& =\rho\left(x_{j}\right)+2\left(g_{j k}^{T} r_{j}\right) \gamma_{j k} \beta_{j k}+d_{j k} \gamma_{j k}^{2} \quad \text { (see (A.4)) }
\end{aligned}
$$

or, equivalently,

$$
2\left(g_{j k}^{T} r_{j}\right) \gamma_{j k} \beta_{j k}+d_{j k} \gamma_{j k}^{2}=\rho\left(x_{k}\right)-\rho\left(x_{j}\right)=\left(\rho\left(x_{k}\right)-\lambda_{1}\right)-\left(\rho\left(x_{j}\right)-\lambda_{1}\right)
$$

Using Assumption 4.3, it follows that

$$
-\left(1-\underline{\xi}^{k-j}\right)\left(\rho\left(x_{j}\right)-\lambda_{1}\right) \leq 2\left(g_{j k}^{T} r_{j}\right) \gamma_{j k} \beta_{j k}+d_{j k} \gamma_{j k}^{2} \leq-\left(1-\bar{\xi}^{k-j}\right)\left(\rho\left(x_{j}\right)-\lambda_{1}\right) .
$$

From (4.4), we divide the above inequality by $\sin ^{2} \theta_{j}$ and have

$$
\begin{align*}
-\left(1-\underline{\xi}^{k-j}\right)\left(\rho\left(f_{j}\right)-\lambda_{1}\right) & \leq 2 \frac{g_{j k}^{T} r_{j}}{\sin \theta_{j}}\left(\frac{\gamma_{j k}}{\sin \theta_{j}}\right) \beta_{j k}+d_{j k}\left(\frac{\gamma_{j k}}{\sin \theta_{j}}\right)^{2}  \tag{A.6}\\
& \leq-\left(1-\bar{\xi}^{k-j}\right)\left(\rho\left(f_{j}\right)-\lambda_{1}\right)
\end{align*}
$$

Note that the left-hand side and the right-hand side of the above inequality are both on the order of $\mathcal{O}(1)$, bounded away from zero and independent of $\theta_{j}$. Also note that $\frac{\left|g_{j r_{j}}^{T} r_{j}\right|}{\sin \theta_{j}} \leq \frac{\left\|g_{j k}\right\|\left\|r_{j}\right\|}{\sin \theta_{j}} \leq(1+\delta)\left\|g_{j k}\right\|\left\|A f_{j}-\rho\left(x_{j}\right) B f_{j}\right\|=\mathcal{O}(1)$ from (4.5). In addition, since $g_{j k} \in \operatorname{span}\left\{p_{j-1}, g_{j}, \ldots, g_{k-1}\right\} \subset W_{k}$, under Assumption 4.4, there is a $d>0$ such that $d_{j k}=\rho\left(g_{j k}\right)-\rho\left(x_{j}\right) \geq d>0$ for all $1 \leq j<k$. Consequently, letting $\theta_{j} \rightarrow 0$ in (A.6), we have $\gamma_{j k}=\mathcal{O}\left(\sin \theta_{j}\right) .{ }^{2}$

From (A.4), $\beta_{j k}$ can be written in terms of $\gamma_{j k}$ as

$$
\begin{equation*}
\beta_{j k}=-\mu_{j k} \gamma_{j k}+\sqrt{1-\gamma_{j k}^{2}+\mu_{j k}^{2} \gamma_{j k}^{2}}=1-\mu_{j k} \gamma_{j k}-\frac{1}{2}\left(1-\mu_{j k}^{2}\right) \gamma_{j k}^{2}+\mathcal{O}\left(\gamma_{j k}^{4}\right) \tag{A.7}
\end{equation*}
$$

which completes the proof.
Proof of Theorem 6. Since $y=\beta_{y} x_{0}+\gamma_{y} g_{y}, g_{y} \in W_{k},\|y\|_{B}=\left\|g_{y}\right\|_{B}=1$, and $\rho(y) \leq \rho\left(x_{0}\right)$, from Lemma 5, we have $\gamma_{y}=\mathcal{O}\left(\sin \theta_{0}\right)$ and $\left|1-\beta_{y}\right|=\mathcal{O}\left(\sin \theta_{0}\right)$. Similarly, the iterate $x_{k}=\beta_{0 k} x_{0}+\gamma_{0 k} g_{0 k}$ satisfies $\gamma_{0 k}=\mathcal{O}\left(\sin \theta_{0}\right)$ and $\left|1-\beta_{0 k}\right|=$ $\mathcal{O}\left(\sin \theta_{0}\right)$. Note that $\left\|r_{0}\right\|=\mathcal{O}\left(\sin \theta_{0}\right), 0 \leq \rho\left(x_{0}\right)-\rho(y) \leq \rho\left(x_{0}\right)-\lambda_{1}=\mathcal{O}\left(\sin ^{2} \theta_{0}\right)$, and, similarly, $0 \leq \rho\left(x_{0}\right)-\rho\left(x_{k}\right) \leq \mathcal{O}\left(\sin ^{2} \theta_{0}\right)$. It follows that

$$
\begin{aligned}
& \text { (A.8) } \quad(A-\rho(y) B) y-\left(A-\rho\left(x_{k}\right) B\right) x_{k} \\
& \begin{array}{l}
=\left(A-\rho\left(x_{0}\right) B\right)\left(\beta_{y} x_{0}+\gamma_{y} g_{y}\right)-\left(A-\rho\left(x_{0}\right) B\right)\left(\beta_{0 k} x_{0}+\gamma_{0 k} g_{0 k}\right) \\
\quad \quad \quad+\left(\rho\left(x_{0}\right)-\rho(y)\right) B y-\left(\rho\left(x_{0}\right)-\rho\left(x_{k}\right)\right) B x_{k} \\
= \\
=\left(\beta_{y}-\beta_{0 k}\right)\left(A-\rho\left(x_{0}\right) B\right) x_{0}+\left(A-\rho\left(x_{0}\right) B\right)\left(\gamma_{y} g_{y}-\gamma_{0 k} g_{0 k}\right)+\mathcal{O}\left(\sin ^{2} \theta_{0}\right) \\
=\left[\left(\beta_{y}-1\right)+\left(1-\beta_{0 k}\right)\right] r_{0}+\left(A-\rho\left(x_{k}\right) B\right)\left(\gamma_{y} g_{y}-\gamma_{0 k} g_{0 k}\right) \\
\quad \quad \quad+\left(\rho\left(x_{k}\right)-\rho\left(x_{0}\right)\right) B\left(\gamma_{y} g_{y}-\gamma_{0 k} g_{0 k}\right)+\mathcal{O}\left(\sin ^{2} \theta_{0}\right) \\
=r_{0} \mathcal{O}\left(\sin \theta_{0}\right)+\gamma_{y}\left(A-\rho\left(x_{k}\right) B\right) g_{y}-\gamma_{0 k}\left(A-\rho\left(x_{k}\right) B\right) g_{0 k}+\mathcal{O}\left(\sin ^{3} \theta_{0}\right)+\mathcal{O}\left(\sin ^{2} \theta_{0}\right) \\
=\gamma_{y}\left(A-\rho\left(x_{k}\right) B\right) g_{y}-\gamma_{0 k}\left(A-\rho\left(x_{k}\right) B\right) g_{0 k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right)
\end{array}
\end{aligned}
$$

By assumption, $g_{y}, g_{0 k} \in W_{k}=\operatorname{span}\left\{p_{0}, \ldots, p_{k-1}\right\}$ are approximately conjugate to $p_{k}$, such that $p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{y}=\mathcal{O}\left(\sin \theta_{0}\right)$ and $p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{0 k}=\mathcal{O}\left(\sin \theta_{0}\right)$, $\operatorname{since} \sin \theta_{\ell}=\mathcal{O}\left(\sin \theta_{0}\right)$ for $1 \leq \ell \leq k-1$, due to (4.9). It follows from (A.8) that

$$
\begin{aligned}
& p_{k}^{T}\left[(A-\rho(y) B) y-\left(A-\rho\left(x_{k}\right) B\right) x_{k}\right] \\
= & \gamma_{y} p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{y}-\gamma_{0 k} p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) g_{0 k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right) \\
= & \gamma_{y} \mathcal{O}\left(\sin \theta_{0}\right)-\gamma_{0 k} \mathcal{O}\left(\sin \theta_{0}\right)+\mathcal{O}\left(\sin ^{2} \theta_{0}\right)=\mathcal{O}\left(\sin \theta_{0}^{2}\right),
\end{aligned}
$$

which is a crucial observation for the rest of the proof, or equivalently,

$$
\begin{equation*}
p_{k}^{T}(A-\rho(y) B) y=p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) x_{k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right) \tag{A.9}
\end{equation*}
$$

In addition, note that $x_{k}-y=\left(\beta_{0 k}-\beta_{y}\right) x_{0}+\left(\gamma_{0 k} g_{0 k}-\gamma_{y} g_{y}\right)=x_{0} \mathcal{O}\left(\sin \theta_{0}\right)+$ $g_{0 k} \mathcal{O}\left(\sin \theta_{0}\right)-g_{y} \mathcal{O}\left(\sin \theta_{0}\right)=\mathcal{O}\left(\sin \theta_{0}\right)$, and hence $B y=B x_{k}+\mathcal{O}\left(\sin \theta_{0}\right)$.

[^2]Let $R_{3}(\rho(x), \alpha p)=\frac{1}{2} \rho(x+\alpha p)-\left[\frac{1}{2} \rho(x)+\alpha p^{T} \nabla \frac{1}{2} \rho(x)+\frac{1}{2} \alpha^{2} p^{T} \nabla^{2} \frac{1}{2} \rho(x) p\right]$ be the remainder of second order Taylor expansion of $\frac{1}{2} \rho(x+\alpha p)$ at $x$ with $\|x\|_{B}=1$. Then,

$$
\begin{aligned}
& \text { (A.10) } \\
& \frac{1}{2} \rho(z)=\frac{1}{2} \rho\left(y+\alpha p_{k}\right)=\frac{1}{2} \rho(y)+\alpha p_{k}^{T} \frac{1}{2} \nabla \rho(y)+\frac{1}{2} \alpha^{2} p_{k}^{T} \frac{1}{2} \nabla^{2} \rho(y) p_{k}+R_{3}\left(\rho(y), \alpha p_{k}\right) \\
& =\frac{1}{2} \rho(y)+\alpha p_{k}^{T}(A-\rho(y) B) y+\frac{1}{2} \alpha^{2} p_{k}^{T}\left[(A-\rho(y) B)-2(A y-\rho(y) B y)(B y)^{T}\right. \\
& \left.-2 B y(A y-\rho(y) B y)^{T}\right] p_{k}+R_{3}\left(\rho(y), \alpha p_{k}\right) \\
& =\frac{1}{2} \rho(y)+\alpha\left[p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) x_{k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right)\right]+\frac{1}{2} \alpha^{2}\left\{p_{k}^{T}\left[\left(A-\rho\left(x_{k}\right) B\right)+\left(\rho\left(x_{k}\right)-\rho(y)\right) B\right] p_{k}\right. \\
& -2\left[p_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) x_{k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right)\right]\left[x_{k}^{T} B+\mathcal{O}\left(\sin \theta_{0}\right)\right] p_{k} \\
& \left.-2 p_{k}^{T}\left[B x_{k}+\mathcal{O}\left(\sin \theta_{0}\right)\right]\left[x_{k}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{k}+\mathcal{O}\left(\sin ^{2} \theta_{0}\right)\right]\right\}+R_{3}\left(\rho(y), \alpha p_{k}\right) \quad(\text { see (A.9)) } \\
& =\frac{1}{2} \rho(y)+\alpha p_{k}^{T} \frac{1}{2} \nabla \rho\left(x_{k}\right)+\alpha \mathcal{O}\left(\sin ^{2} \theta_{0}\right)+\frac{1}{2} \alpha^{2} p_{k}^{T}\left\{\left(A-\rho\left(x_{k}\right) B\right)-2 r_{k}\left(B x_{k}\right)^{T}-2\left(B x_{k}\right) r_{k}^{T}\right\} p_{k} \\
& +\alpha^{2}\left[\mathcal{O}\left(\sin ^{2} \theta_{0}\right)+\left\|r_{k}\right\| \mathcal{O}\left(\sin \theta_{0}\right)\right]+R_{3}\left(\rho(y), \alpha p_{k}\right) \\
& =\frac{1}{2} \rho(y)+\alpha p_{k}^{T} \frac{1}{2} \nabla \rho\left(x_{k}\right)+\frac{1}{2} \alpha^{2} p_{k}^{T} \frac{1}{2} \nabla^{2} \rho\left(x_{k}\right) p_{k}+\alpha \mathcal{O}\left(\sin ^{2} \theta_{0}\right)+\alpha^{2} \mathcal{O}\left(\sin ^{2} \theta_{0}\right)+R_{3}\left(\rho(y), \alpha p_{k}\right) \\
& =\frac{1}{2} \rho\left(x_{k}\right)+\alpha p_{k}^{T} \frac{1}{2} \nabla \rho\left(x_{k}\right)+\frac{1}{2} \alpha^{2} p_{k}^{T} \frac{1}{2} \nabla^{2} \rho\left(x_{k}\right) p_{k}+R_{3}\left(\rho\left(x_{k}\right), \alpha p_{k}\right) \\
& +\frac{1}{2}\left(\rho(y)-\rho\left(x_{k}\right)\right)-R_{3}\left(\rho\left(x_{k}\right), \alpha p_{k}\right)+\mathcal{O}\left(\alpha \sin ^{2} \theta_{0}\right)+\mathcal{O}\left(\alpha^{2} \sin ^{2} \theta_{0}\right)+R_{3}\left(\rho(y), \alpha p_{k}\right) \\
& =\frac{1}{2} \rho\left(x_{k}+\alpha p_{k}\right)+\frac{1}{2}\left(\rho(y)-\rho\left(x_{k}\right)\right)+R_{3}\left(\rho(y), \alpha p_{k}\right)-R_{3}\left(\rho\left(x_{k}\right), \alpha p_{k}\right)+\mathcal{O}\left(\alpha \sin ^{2} \theta_{0}\right)+\mathcal{O}\left(\alpha^{2} \sin ^{2} \theta_{0}\right) .
\end{aligned}
$$

Let the global minimizer in $U_{k+1}$ be $z^{*}=y\left(z^{*}\right)+\alpha\left(z^{*}\right) p_{k}$ with $y\left(z^{*}\right) \in U_{k}$ and $\left\|y\left(z^{*}\right)\right\|_{B}=1$. We note that $y\left(z^{*}\right)$ is generally not the global minimizer in $U_{k}$. Consider the decomposition $y\left(z^{*}\right)=v_{1} \cos \theta_{y\left(z^{*}\right)}+f_{y\left(z^{*}\right)} \sin \theta_{y\left(z^{*}\right)}$ with $f_{y\left(z^{*}\right)} \perp B v_{1}$ and $\left\|f_{y\left(z^{*}\right)}\right\|_{B}=1$ such that $\rho\left(y\left(z^{*}\right)\right)-\lambda_{1}=\mathcal{O}\left(\sin ^{2} \theta_{y\left(z^{*}\right)}\right)$. Here, $\alpha\left(z^{*}\right)$ is the optimal step size moving from $y\left(z^{*}\right)$ in the direction of $p_{k}$, due to the global optimality of $z^{*}$ in $U_{k+1}$. It follows from Lemma 4 that

$$
\rho\left(z^{*}\right)-\rho\left(y\left(z^{*}\right)\right) \leq-\mathcal{O}\left(\sin ^{2} \theta_{y\left(z^{*}\right)}\right) \mathcal{O}\left(\cos ^{2} \angle\left(p_{k}, \nabla \rho\left(y\left(z^{*}\right)\right)\right)\right)
$$

where the coefficient of the $\sin ^{2} \theta_{y\left(z^{*}\right)}$ term depends on the quantities $b$ and $c$ defined in (4.12) involving the search direction $p_{k}$. On the other hand, as $z^{*}$ is the global minimizer in $U_{k+1}$, we have

$$
\begin{equation*}
\rho\left(z^{*}\right)-\lambda_{1} \leq \rho\left(x_{k+1}\right)-\lambda_{1}=\mathcal{O}\left(\sin ^{2} \theta_{k+1}\right) \tag{A.11}
\end{equation*}
$$

Meanwhile, by Lemma 4, we also have

$$
\text { 12) } \begin{align*}
& \rho\left(z^{*}\right)-\lambda_{1}=\left[\rho\left(z^{*}\right)-\rho\left(y\left(z^{*}\right)\right)\right]+\left[\rho\left(y\left(z^{*}\right)\right)-\lambda_{1}\right]  \tag{A.12}\\
=- & -\mathcal{O}\left(\sin ^{2} \theta_{y\left(z^{*}\right)}\right) \mathcal{O}\left(\cos ^{2} \angle\left(p_{k}, \nabla \rho\left(y\left(z^{*}\right)\right)\right)\right)+\mathcal{O}\left(\sin ^{2} \theta_{y\left(z^{*}\right)}\right)=\mathcal{O}\left(\sin ^{2} \theta_{y\left(z^{*}\right)}\right) .
\end{align*}
$$

It follows that

$$
\begin{equation*}
\sin \theta_{y\left(z^{*}\right)} \leq \mathcal{O}\left(\sin \theta_{k+1}\right) \tag{A.13}
\end{equation*}
$$

Given $z^{*}=y\left(z^{*}\right)+\alpha\left(z^{*}\right) p_{k}$, it follows from the triangle inequality that

$$
\begin{equation*}
\left\|z^{*}\right\|_{B} \leq\left\|y\left(z^{*}\right)\right\|_{B}+\left|\alpha\left(z^{*}\right)\right|\left\|p_{k}\right\|_{B}=1+\left|\alpha\left(z^{*}\right)\right| \tag{A.14}
\end{equation*}
$$

Meanwhile, the $B$-normalized $z^{*}$ is $\frac{z^{*}}{\left\|z^{*}\right\|_{B}}=\frac{1}{\left\|z^{*}\right\|_{B}} y\left(z^{*}\right)+\frac{\alpha\left(z^{*}\right)}{\left\|z^{*}\right\|_{B}} p_{k}$, and we can follow the proof of Lemma 5 to show that $\frac{\left|\alpha\left(z^{*}\right)\right|}{\left\|z^{*}\right\|_{B}}=\mathcal{O}\left(\sin \theta_{y\left(z^{*}\right)}\right) \leq \mathcal{O}\left(\sin \theta_{k+1}\right)$ due to (A.13).
From (A.14), $\frac{\left|\alpha\left(z^{*}\right)\right|}{1+\left|\alpha\left(z^{*}\right)\right|} \leq \frac{\left|\alpha\left(z^{*}\right)\right|}{\left\|z^{*}\right\|_{B}} \leq \mathcal{O}\left(\sin \theta_{k+1}\right)$, i.e., $\left|\alpha\left(z^{*}\right)\right| \leq \mathcal{O}\left(\sin \theta_{k+1}\right)$.

Let $\alpha_{k}^{*}$ be the minimizer of $\rho\left(x_{k}+\alpha p_{k}\right)$ and $y^{*}$ be the global minimizer in $U_{k}$ (note that $\alpha_{k}^{*} \neq \alpha\left(z^{*}\right)$ and $y^{*} \neq y\left(z^{*}\right)$ ), which contains all vectors of the form $y=$ $\beta_{y} x_{0}+\gamma_{y} g_{y}$ with $\|y\|_{B}=\left\|g_{y}\right\|_{B}=1, g_{y} \in W_{k}$, and $\rho(y) \leq \rho\left(x_{0}\right)$. Let $z=z^{*}=$ $y\left(z^{*}\right)+\alpha\left(z^{*}\right) p_{k}, y=y\left(z^{*}\right)$, and $\alpha=\alpha\left(z^{*}\right)$ in (A.10), and note that $\rho\left(x_{k}+\alpha\left(z^{*}\right) p_{k}\right) \geq$ $\rho\left(x_{k}+\alpha^{*} p_{k}\right)$ and $\rho\left(y\left(z^{*}\right)\right) \geq \rho\left(y^{*}\right)$. Therefore, we have

$$
\begin{equation*}
\frac{1}{2} \rho\left(z^{*}\right)=\frac{1}{2} \rho\left(x_{k}+\alpha\left(z^{*}\right) p_{k}\right)+\frac{1}{2}\left(\rho\left(y\left(z^{*}\right)\right)-\rho\left(x_{k}\right)\right)+R_{3}\left(\rho\left(y\left(z^{*}\right)\right), \alpha\left(z^{*}\right) p_{k}\right) \tag{A.15}
\end{equation*}
$$

$$
-R_{3}\left(\rho\left(x_{k}\right), \alpha\left(z^{*}\right) p_{k}\right)+\mathcal{O}\left(\alpha\left(z^{*}\right) \sin ^{2} \theta_{0}\right)+\mathcal{O}\left(\alpha\left(z^{*}\right)^{2} \sin ^{2} \theta_{0}\right)
$$

$$
\geq \frac{1}{2} \rho\left(x_{k}+\alpha_{k}^{*} p_{k}\right)+\frac{1}{2}\left(\rho\left(y^{*}\right)-\rho\left(x_{k}\right)\right)+R_{3}\left(\rho\left(y\left(z^{*}\right)\right), \alpha\left(z^{*}\right) p_{k}\right)
$$

$$
-R_{3}\left(\rho\left(x_{k}\right), \alpha\left(z^{*}\right) p_{k}\right)+\mathcal{O}\left(\alpha\left(z^{*}\right) \sin ^{2} \theta_{0}\right)+\mathcal{O}\left(\alpha\left(z^{*}\right)^{2} \sin ^{2} \theta_{0}\right)
$$

$$
\geq \frac{1}{2} \rho\left(x_{k+1}\right)+\frac{1}{2}\left(\rho\left(y^{*}\right)-\rho\left(x_{k}\right)\right)+\mathcal{O}\left(\sin \theta_{k+1}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right) . \quad\left(\left|\alpha\left(z^{*}\right)\right| \leq \mathcal{O}\left(\sin \theta_{k+1}\right)\right)
$$

In the last step above, $\rho\left(x_{k}+\alpha^{*} p_{k}\right) \geq \rho\left(x_{k+1}\right)$ because $x_{k+1}$ is the minimizer over $\operatorname{span}\left\{x_{k}, g_{k}, p_{k-1}\right\}$, whereas $p_{k} \in \operatorname{span}\left\{g_{k}, p_{k-1}\right\}$. Equivalently,

$$
\rho\left(x_{k+1}\right)-\rho\left(z^{*}\right) \leq \rho\left(x_{k}\right)-\rho\left(y^{*}\right)+\mathcal{O}\left(\sin \theta_{k+1}\right) \mathcal{O}\left(\sin ^{2} \theta_{0}\right)
$$

Proof of Lemma 7. Recall that $\widetilde{p}_{k}=g_{k}-\frac{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}} p_{k-1}$. Therefore

$$
\begin{aligned}
& p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) \widetilde{p}_{k} \\
= & p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}-\frac{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) g_{k}}{p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}} p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k-1}=0,
\end{aligned}
$$

and hence $p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k}=0$. Note that $\rho\left(x_{k-1}\right)-\rho\left(x_{k}\right)=\left(\rho\left(x_{k-1}\right)-\lambda_{1}\right)-$ $\left(\rho\left(x_{k}\right)-\lambda_{1}\right) \leq(1-\bar{\xi})\left(\rho\left(x_{k-1}\right)-\lambda_{1}\right)=\mathcal{O}\left(\sin ^{2} \theta_{k-1}\right)$, and it follows that

$$
\begin{aligned}
& p_{k-1}^{T}\left(A-\rho\left(x_{k}\right) B\right) p_{k}=p_{k-1}^{T}\left(A-\rho\left(x_{k-1}\right) B\right) p_{k}+\left(\rho\left(x_{k-1}\right)-\rho\left(x_{k}\right)\right) p_{k-1}^{T} B p_{k} \\
= & \left(\rho\left(x_{k-1}\right)-\rho\left(x_{k}\right)\right)\left(p_{k-1}, p_{k}\right)_{B}=\mathcal{O}\left(\sin ^{2} \theta_{k-1}\right)
\end{aligned}
$$

Proof of Lemma 8. At step 2 of Algorithm 2, $x_{2}$ is extracted from $\operatorname{span}\left\{x_{1}, g_{1}, p_{0}\right\}$, where $p_{0}=g_{0}=M d_{0}$ up to a scaling factor. By the local optimality,

$$
\begin{equation*}
r_{2} \perp g_{0} \text { and } r_{2} \perp g_{1} \tag{A.16}
\end{equation*}
$$

At step 3 , we form $\operatorname{span}\left\{x_{2}, g_{2}, p_{1}\right\}$ to extract $x_{3}$ and $p_{2}$. At step 7 of Algorithm $2, \widetilde{p}_{2}=$ $g_{2}-\frac{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) g_{2}}{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) p_{1}} p_{1}$, which shows that $\left\|\widetilde{p}_{2}\right\|$ is proportional to $\left\|g_{2}\right\|$ and independent of the scaling of $p_{1}$. Hence, the normalized search direction $p_{2}$ can be written as $p_{2}=\frac{\eta_{2}}{\left\|g_{2}\right\|} \widetilde{p}_{2}$, where $\eta_{2}$ is chosen such that $\left\|p_{2}\right\|_{B}=1$. From Lemma 5 , we have $x_{1}=\beta_{01} x_{0}+\gamma_{01} p_{0}$, i.e., $p_{0}=\frac{1}{\gamma_{01}}\left(-\beta_{01} x_{0}+x_{1}\right)$, where $\gamma_{01}=\mathcal{O}\left(\sin \theta_{0}\right)$. Therefore

$$
\begin{aligned}
& p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2}=\frac{1}{\gamma_{01}}\left[-\beta_{01} x_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right)+x_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right)+\left(\rho\left(x_{1}\right)-\rho\left(x_{0}\right)\right) x_{1}^{T} B\right] g_{2} \\
& =\frac{\rho\left(x_{1}\right)-\rho\left(x_{0}\right)}{\gamma_{01}} x_{1}^{T} B g_{2}+\frac{1}{\gamma_{01}}\left(-\beta_{01} r_{0}^{T}+r_{1}^{T}\right) g_{2}=\frac{\left(\rho\left(x_{1}\right)-\lambda_{1}\right)-\left(\rho\left(x_{0}\right)-\lambda_{1}\right)}{\gamma_{01}} x_{1}^{T} B g_{2} \\
& +\frac{1}{\gamma_{01}}\left(-\beta_{01} r_{0}^{T} M p_{m-1}^{(2)}\left(\left(A-\rho_{2} B\right) M\right) r_{2}+r_{1}^{T} M p_{m-1}^{(2)}\left(\left(A-\rho_{2} B\right) M\right) r_{2}\right) \\
& =\frac{\mathcal{O}\left(\sin ^{2} \theta_{0}\right)}{\mathcal{O}\left(\sin \theta_{0}\right)} x_{1}^{T} B g_{2}+\frac{1}{\gamma_{01}}\left\{-\beta_{01} r_{0}^{T}\left(p_{m-1}^{(0)}\left(M\left(A-\rho_{0} B\right)\right) M+\mathcal{O}\left(\sin \theta_{0}\right)\right) r_{2}\right. \\
& \left.\quad+r_{1}^{T}\left(p_{m-1}^{(1)}\left(M\left(A-\rho_{1} B\right)\right) M+\mathcal{O}\left(\sin \theta_{1}\right)\right) r_{2}\right\} .
\end{aligned}
$$

Since $M, A$, and $B$ are all real symmetric, so is $p_{m-1}^{(k)}\left(M\left(A-\rho_{k} B\right)\right) M$. Therefore, $r_{k}^{T} p_{m-1}^{(k)}\left(M\left(A-\rho_{k} B\right)\right) M=g_{k}^{T}$, and hence

$$
\begin{aligned}
& p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2} \\
= & x_{1}^{T} B g_{2} \mathcal{O}\left(\sin \theta_{0}\right)+\frac{1}{\gamma_{01}}\left\{-\beta_{01} g_{0}^{T} r_{2}+r_{0}^{T} r_{2} \mathcal{O}\left(\sin \theta_{0}\right)+g_{1}^{T} r_{2}+r_{1}^{T} r_{2} \mathcal{O}\left(\sin \theta_{1}\right)\right\} \\
= & \left\|g_{2}\right\| \mathcal{O}\left(\sin \theta_{0}\right)+\frac{1}{\mathcal{O}\left(\sin \theta_{0}\right)}\left\{r_{0}^{T} r_{2} \mathcal{O}\left(\sin \theta_{0}\right)+r_{1}^{T} r_{2} \mathcal{O}\left(\sin \theta_{1}\right)\right\} \\
= & \left\|g_{2}\right\| \mathcal{O}\left(\sin \theta_{0}\right)+\left\|r_{2}\right\| \mathcal{O}\left(\sin \theta_{0}\right), \quad\left(\text { note that }\left\|r_{1}\right\|=\mathcal{O}\left(\sin \theta_{1}\right)=\mathcal{O}\left(\sin \theta_{0}\right)\right),
\end{aligned}
$$

where we use $r_{2} \perp g_{0}$ and $r_{2} \perp g_{1}$ from (A.16). This shows that $p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2} \leq$ $\mathcal{O}\left(\sin \theta_{0}\right)$. Also, since $p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2}$ is proportional to $\left\|g_{2}\right\|$, we simply have $p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2}=\left\|g_{2}\right\| \mathcal{O}\left(\sin \theta_{0}\right)$. Therefore, by Lemma 7 ,

$$
\begin{aligned}
& p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) p_{2}=p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) \frac{\eta_{2}}{\left\|g_{2}\right\|} \widetilde{p}_{2} \\
= & p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) \frac{\eta_{2}}{\left\|g_{2}\right\|}\left\{g_{2}-\frac{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) g_{2}}{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) p_{1}} p_{1}\right\} \\
= & \frac{\eta_{2}}{\left\|g_{2}\right\|}\left\{p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) g_{2}-\frac{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) g_{2}}{p_{1}^{T}\left(A-\rho\left(x_{1}\right) B\right) p_{1}}\left[p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) p_{1}\right]\right\} \\
= & \frac{\eta_{2}}{\left\|g_{2}\right\|}\left(\left\|g_{2}\right\| \mathcal{O}\left(\sin \theta_{0}\right)+\left\|g_{2}\right\| \mathcal{O}\left(\sin ^{2} \theta_{0}\right)\right)=\mathcal{O}\left(\sin \theta_{0}\right) .
\end{aligned}
$$

Similarly, using a slightly different shift in the above relation, we have

$$
\begin{align*}
& p_{2}^{T}\left(A-\rho\left(x_{2}\right) B\right) p_{0}=p_{0}^{T}\left(A-\rho\left(x_{0}\right) B\right) p_{2}+\left(\rho\left(x_{0}\right)-\rho\left(x_{2}\right)\right) p_{0}^{T} B p_{2}  \tag{A.17}\\
= & \mathcal{O}\left(\sin \theta_{0}\right)+\left(\left(\rho\left(x_{0}\right)-\lambda_{1}\right)-\left(\rho\left(x_{2}\right)-\lambda_{1}\right)\right)\left(p_{0}, p_{2}\right)_{B} \\
= & \mathcal{O}\left(\sin \theta_{0}\right)+\left(\mathcal{O}\left(\sin ^{2} \theta_{0}\right)-\mathcal{O}\left(\sin ^{2} \theta_{2}\right)\right)\left(p_{0}, p_{2}\right)_{B}=\mathcal{O}\left(\sin \theta_{0}\right)
\end{align*}
$$

Proof of Lemma 9. The proof is done by contradiction. Assume that $\lim _{\theta_{0} \rightarrow 0} \cos$ $\angle\left(r_{k}, W_{k}\right)$ exists but is not zero. Then there is a vector $u \in U_{k}$ such that $\lim _{\theta_{0} \rightarrow 0} \cos$ $\angle\left(r_{k}, u\right)=\delta<0$, i.e., $u$ is a descent direction at $x_{k}$. We have shown in Lemma 4 that

$$
\rho\left(x_{k}+\alpha^{*} u\right)-\rho\left(x_{k}\right) \leq-\mathcal{O}\left(\sin ^{2} \theta_{k}\right) \mathcal{O}\left(\cos ^{2} \angle\left(u, r_{k}\right)\right)=-\mathcal{O}\left(\sin ^{2} \theta_{k}\right)
$$

Since $y^{*}$ is the global minimizer in $U_{k}$ and $x_{k}+\alpha^{*} u \in U_{k}, \rho\left(y^{*}\right) \leq \rho\left(x_{k}+\alpha^{*} u\right) \leq \rho\left(x_{k}\right)$. As a result, we have

$$
\lim _{\theta_{0} \rightarrow 0} \frac{\rho\left(x_{k}\right)-\rho\left(y^{*}\right)}{\rho\left(x_{k}\right)-\lambda_{1}} \geq \lim _{\theta_{0} \rightarrow 0} \frac{\rho\left(x_{k}\right)-\rho\left(x_{k}+\alpha^{*} u\right)}{\rho\left(x_{k}\right)-\lambda_{1}} \geq \lim _{\theta_{0} \rightarrow 0} \frac{\mathcal{O}\left(\sin ^{2} \theta_{k}\right)}{\mathcal{O}\left(\sin ^{2} \theta_{k}\right)}=\mathcal{O}(1)
$$

meaning that $x_{k}$ does not satisfy the global quasi optimality condition.
Acknowledgment. The authors are indebted to the referees for their meticulous reading and constructive comments.

## REFERENCES

[1] J. Baglama and L. Reichel, Augmented implicitly restarted Lanczos bidiagonalization methods, SIAM J. Sci. Comput., 27 (2005), pp. 19-42.
[2] D. Calvetti, L. Reichel, and D. C. Sorensen, An implicitly restarted Lanczos method for large symmetric eigenvalue problems, Electron. Trans. Numer. Anal., 2 (1994), p. 21.
[3] A. Chapman and Y. SaAd, Deflated and augmented Krylov subspace techniques, Numer. Linear Algebra Appl., 4 (1997), pp. 43-66.
[4] J. Chen, L. Wu, K. Audhkhasi, B. Kingsbury, and B. Ramabhadrari, Efficient one-vs.-one kernel ridge regression for speech recognition, in Proceedings of the 2016 IEEE International Conference on Acoustics, Speech and Signal Processing, 2016, pp. 2454-2458.
[5] M. Crouzeix, B. Philippe, and M. Sadkane, The Davidson method, SIAM J. Sci. Comput., 15 (1994), pp. 62-76.
[6] E. R. DaVIDSon, The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices, J. Comput. Phys., 17 (1975), pp. 87-94.
[7] E. de Sturler, Nested Krylov methods based on GCR, J. Comput. Appl. Math., 67 (1996), pp. 15-41.
[8] E. DE Sturler, Truncation strategies for optimal Krylov subspace methods, SIAM J. Matrix Anal. Appl., 36 (1999), pp. 864-889.
[9] A. Edelman, T. A. Arias, and S. T. Smith, The geometry of algorithms with orthogonality constraints, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 303-353.
[10] G. H. Golub and C. F. Van Loan, Matrix Computations, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
[11] G. H. Golub and Q. Ye, An inverse free preconditioned Krylov subspace method for symmetric generalized eigenvalue problems, SIAM J. Sci. Comput., 24 (2002), pp. 312-334.
[12] C. Greif, S. He, and P. Liu, SYM-ILDL: Incomplete LDL ${ }^{\top}$ Factorization of Symmetric Indefinite and Skew-Symmetric Matrices, CoRR, arXiv:1505.07589, 2015.
[13] J. Han, M. Kamber, and J. Pei, Data Mining: Concepts and Techniques, 3rd ed., Morgan Kaufmann, San Francisco, CA, 2011.
[14] V. Hernandez, J. Roman, A. Tomas, and V. Vidal, Lanczos Methods in SLEPc, SLEPc Technical Report STR-5, Universidad Politécnica de Valencia, Valencia, Spain, 2006.
[15] C.-J. Hsieh and P. Olsen, Nuclear norm minimization via active subspace selection, in Proceedings of the 31st International Conference on Machine Learning, 2014, pp. 575-583.
[16] V. Kalantzis, R. Li, and Y. SaAd, Spectral Schur complement techniques for symmetric eigenvalue problems, Electron. Trans. Numer. Anal., 45 (2016), pp. 305-329.
[17] A. V. Knyazev, Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method, SIAM J. Sci. Comput., 23 (2001), pp. 517-541.
[18] A. V. Knyazev, M. E. Argentati, I. Lashuk, and E. E. Ovtchinnikov, Block locally optimal preconditioned eigenvalue xolvers (BLOPEX) in HYPRE and PETSC, SIAM J. Sci. Comput., 29 (2007), pp. 2224-2239.
[19] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, J. Res. Nat. Bureau Standards Sec. B, 45 (1950), pp. 255-282.
[20] R. Li, Y. Xi, E. Vecharynski, C. Yang, and Y. Saad, A thick-restart Lanczos algorithm with polynomial filtering for Hermitian eigenvalue problems, SIAM J. Sci. Comput., 38 (2016), pp. A2512-A2534.
[21] X. Liu, Z. Wen, and Y. Zhang, Limited memory block Krylov subspace optimization for computing dominant singular value decompositions, SIAM J. Sci. Comput., 35 (2013), pp. A1641-A1668.
[22] R. Morgan, On restarting the Arnoldi method for large nonsymmetric eigenvalue problems, Math. Comp., 65 (1996), pp. 1213-1230.
[23] R. B. Morgan and D. S. Scott, Generalizations of Davidsons method for computing eigenvalues of sparse symmetric matrices, SIAM J. Sci. Comput., 7 (1986), pp. 817-825.
[24] R. B. Morgan and D. S. Scott, Preconditioning the Lanczos algorithm for sparse symmetric eigenvalue problems, SIAM J. Sci. Comput., 14 (1993), pp. 585-593.
[25] C. W. Murray, S. C. Racine, and E. R. Davidson, Improved algorithms for the lowest few eigenvalues and associated eigenvectors of large matrices, J. Comput. Phys., 103 (1992), pp. 382-389.
[26] J. Olsen, P. Jørgensen, and J. Simons, Passing the one-billion limit in full configurationinteraction (FCI) calculations, Chem. Phys. Lett., 169 (1990), pp. 463-472.
[27] C. C. Paige, Computational variants of the Lanczos method for the eigenproblem, IMA J. Appl. Math., 10 (1972), pp. 373-381.
[28] C. C. Paige, Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix, IMA J. Appl. Math., 18 (1976), pp. 341-349.
[29] C. C. Paige, Accuracy and effectiveness of the Lanczos algorithm for the symmetric eigenproblem, Linear Algebra Appl., 34 (1980), pp. 235-258.
[30] B. N. Parlett, The Symmetric Eigenvalue Problem, Prentice-Hall, Englewood Cliffs, NJ, 1980.
[31] Y. SaAd, Numerical Methods for Large Eigenvalue Problems, Manchester University Press, Manchester, UK, 1992.
[32] G. L. Sleijpen and H. A. Van der Vorst, A Jacobi-Davidson iteration method for linear eigenvalue problems, SIAM Rev., 42 (2000), pp. 267-293.
[33] D. C. Sorensen, Implicit application of polynomial filters in a $k$-step Arnoldi method, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 357-385.
[34] A. Stathopoulos, Nearly optimal preconditioned methods for Hermitian eigenproblems under limited memory. Part I: Seeking one eigenvalue, SIAM J. Sci. Comput., 29 (2007), pp. 481-514.
[35] A. Stathopoulos and C. F. Fischer, A Davidson program for finding a few selected extreme eigenpairs of a large, sparse, real, symmetric matrix, Comput. Phys. Commun., 79 (1994), pp. 268-290.
[36] A. Stathopoulos and Y. SaAd, Restarting techniques for the (Jacobi-) Davidson symmetric eigenvalue methods, Electron. Trans. Numer. Anal., 7 (1998), pp. 163-181.
[37] A. Stathopoulos, Y. SaAd, and K. Wu, Dynamic thick restarting of the Davidson, and the implicitly restarted Arnoldi methods, SIAM J. Sci. Comput., 19 (1998), pp. 227-245.
[38] D. B. Szyld and F. Xue, Preconditioned eigensolvers for large-scale nonlinear Hermitian eigenproblems with variational characterizations. I. Extreme eigenvalues, Math. Comp., 85 (2016), pp. 2887-2918, https://doi.org/10.1090/mcom/3083.
[39] H. A. VAN DER Vorst, Computational methods for large eigenvalue problems, in Handbook of Numerical Analysis, Handb. Numer. Anal. 8, North-Holland, Amsterdam, 2002, pp. 3-179.
[40] E. Vecharynski, C. Yang, and J. E. Pask, A projected preconditioned conjugate gradient algorithm for computing many extreme eigenpairs of a Hermitian matrix, J. Comput. Phys., 290 (2015), pp. 73-89.
[41] E. Vecharynski, C. Yang, and F. Xue, Generalized preconditioned locally harmonic residual method for non-Hermitian eigenproblems, SIAM J. Sci. Comput., 38 (2016), pp. A500A527, https://doi.org/10.1137/15M1027413.
[42] C. Vömel, S. Z. Tomov, O. A. Marques, A. Canning, L.-W. Wang, and J. J. DonGARRA, State-of-the-art eigensolvers for electronic structure calculations of large scale nano-systems, J. Comput. Phys., 227 (2008), pp. 7113-7124.
[43] K. Wu and H. Simon, Thick-restart Lanczos method for large symmetric eigenvalue problems, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 602-616.
[44] L. Wu, P.-Y. Chen, I. E.-H. Yen, F. Xu, Y. Xia, and C. Aggarwal, Scalable spectral clustering using random binning features, in Proceedings of the 24nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2018, pp. 2506-2515.
[45] L. Wu, J. Laeuchli, V. Kalantzis, A. Stathopoulos, and E. Gallopoulos, Estimating the trace of the matrix inverse by interpolating from the diagonal of an approximate inverse, J. Comput. Phys., 326 (2016), pp. 828-844.
[46] L. Wu, E. Romero, and A. Stathopoulos, PRIMME_SVDS: A high-performance preconditioned svd solver for accurate large-scale computations, SIAM J. Sci. Comput., 39 (2017), pp. S248-S271.
[47] L. Wu and A. Stathopoulos, A preconditioned hybrid SVD method for accurately computing singular triplets of large matrices, SIAM J. Sci. Comput., 37 (2015), pp. S365-S388.
[48] L. Wu, K. J. Wu, A. Sim, M. Churchill, J. Y. Choi, A. Stathopoulos, C.-S. Chang, and S. Klasky, Towards real-time detection and tracking of spatio-temporal features: Blobfilaments in fusion plasma, IEEE Trans. Big Data, 2 (2016), pp. 262-275.
[49] L. Wu, I. E. Yen, J. Chen, and R. Yan, Revisiting random binning features: Fast convergence and strong parallelizability, in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1265-1274.
[50] C. Yang, Solving large-scale eigenvalue problems in SciDAC applications, J. Phys. Conf. Ser., 16 (2005), pp. 425-434.


[^0]:    *Submitted to the journal's Methods and Algorithms for Scientific Computing section November 20, 2017; accepted for publication (in revised form) January 7, 2019; published electronically April 9, 2019.
    http://www.siam.org/journals/sisc/41-2/M115756.html
    Funding: This work was supported by NSF under grants ACI SI2-SSE 1440700 and DMS1719461 and by DOE under grant DE-FC02-12ER41890.
    ${ }^{\dagger}$ Department of Computer Science, College of William and Mary, Williamsburg, VA 23187-8795 (lwu@email.wm.edu, andreas@cs.wm.edu).
    $\ddagger$ Department of Mathematical Sciences, Clemson University, Clemson, SC 29631 (fxue@clemson. edu).

[^1]:    ${ }^{1}$ One could use search directions instead of previous Ritz vectors for better numerical stability as suggested in [17, 18], but both variants essentially construct the same subspace.

[^2]:    ${ }^{2}$ We emphasize here that $\gamma_{j k} \neq o\left(\sin \theta_{j}\right)$ and hence $\gamma_{j k}$ can be safely replaced with $\mathcal{O}\left(\sin \theta_{j}\right)$ anywhere appropriate.

